

Speaking of Knowing

PATRICK RYSIEW

THE UNIVERSITY OF VICTORIA

1. INTRODUCTION

What do we talk about when we talk about knowing? No doubt, when a speaker utters a sentence of the form, ‘*S* knows [or does not know] that *p*,’ the sentence itself expresses the proposition it wears on its sleeve -- that *S* knows [/doesn’t know] that *p*. But is it in order to communicate just that information that a speaker typically utters such a sentence? If so, a strategy suggests itself for doing epistemology: since our everyday knowledge-attributing practices are guided simply by knowledge sentences’ semantic contents, we can look to facts about the circumstances in which we ordinarily attribute knowledge to get a clearer view of the meaning of ‘knows’, the truth conditions of the relevant sentences, perhaps even the nature of knowledge itself.

Precisely this strategy has underwritten a lot of recent epistemological theorizing, and a lot of very influential epistemology: a central source of data for epistemologists remains our everyday knowledge-attributing practices – what we are/aren’t prepared to say about whether a subject knows that *p*, the circumstances under which we typically attribute knowledge, the oddity of certain utterances containing epistemic terms, and so on. Inspired by such data, a number of epistemologists have advanced claims that cast doubt on one or another aspect of our ordinary (‘intuitive’) ways of thinking about knowledge. Some of these claims are of the sceptical variety: contrary to our everyday epistemic pretensions, we’re told, we have hardly any knowledge at all; for as our own knowledge-attributing practices reveal, knowledge is much more demanding than we typically realize – it requires that there be no possibility of error, or that one be entitled to be absolutely certain that one is correct. Other claims are non-, even anti-, sceptical in their intended import. To take a prominent example, a number of recent theorists have suggested that, contrary to what is generally supposed, the conditions under which one can be truly said to know can change with shifts

in ‘context’. As with their sceptical counterparts, such revisionist claims are typically rooted in reflections on our ordinary knowledge-attributing practices. Indeed, it is sometimes implied (e.g., DeRose 2005; Ludlow 2005) that it is a distinctive virtue of such views that they are so based.

Such revisionist arguments aside, and as a quite general point, it isn’t surprising that epistemologists should take our epistemic discourse as straightforwardly revealing of the relevant epistemic phenomena. Because so much of what goes on in linguistic communication is inexplicit, it’s easy to fail to take note of it and, as a result, to suppose that in our talk language serves merely as a sort of transparent vehicle for disclosing our thoughts. But where semantic content is not a reliable guide either to what the speaker is trying to communicate, and so why he says what he does, or to our own reactions to this or that attribution or denial of knowledge, in making inferences from the latter data to the meanings of certain central epistemic terms or concepts, or to the truth conditions of sentences involving them, we run the risk of mistaking pragmatic phenomena for semantic ones.¹

Of course, that’s a point that everyone is liable to grant: epistemologists need to be on guard against letting merely pragmatic factors determine the shape of the relevant semantic theories. Still, most would also insist, along with Lawrence Bonjour, that our intuitive judgements about particular cases, including our intuitive response to various specific attributions and denials of knowledge, “are a central and essential part of our basis for understanding and delineating general concepts such as our concept of knowledge,” and that “if all such judgments were dismissed as undependable, we would have little handle left on such concepts” (Bonjour 2002, p. 32). Agreed. But while our everyday knowledge-attributing practices and intuitive epistemic judgments as a whole may constitute a central, even indispensable, source of evidence for epistemologists, that evidence is defeasible and may be over-ridden in particular cases: when we have good reason to suppose that a specific part or aspect of those practices, or a certain subclass of those judgments, are crucially affected by pragmatic factors and may indeed be the product thereof, we can rightly refrain from letting those data determine the shape of the relevant epistemological theories.

The goal of the present paper is to make a case for the claim that this last, in fact, is the situation

with regard to the revisionist claims mentioned above: Section 2 describes how assertions in general make available more information than what is asserted, including information about the speaker's own epistemic commitments. In Section 3, this general idea is applied to explicit attributions/denials of knowledge, in both their first- and third-personal forms: given some extremely uncontroversial assumptions about knowledge (which together comprise what I call 'the ho-hum view') and a widely-held view of how rational communication works, knowledge ascriptions can be expected to convey much more than that the ho-hum conditions on knowing are satisfied; though still epistemic, such information is stronger and more specific than those uncontroversial assumptions require. Section 4 applies these results to arguments for the revisionist claims mentioned above, arguing that the allegedly pro-revisionist data can be predicted using the ho-hum view plus the pragmatics outlined in Sections 2 and 3. Such data may or may not be reflected in the relevant semantic and/or epistemological facts; but whether or not they are, we can expect such information to be carried by the relevant assertions. Hence, that such information is so carried does not constitute evidence for the revisionist views in question. Further, it will emerge (Section 5) that, in addition to enabling us to undercut certain arguments against the naïve ('ho-hum') conception of knowledge, taking seriously the influence of pragmatic factors on our talk of knowing makes it much harder to be a revisionist. For whereas the naïve conception of knowing plus an agreed-upon pragmatics enables one to predict the allegedly revisionist-friendly data, if one rejects the idea that such pragmatic factors are importantly responsible for that data, the thought that they are friendly to the revisionist becomes difficult to sustain. But while the ho-hum view is motivated by very familiar theoretical considerations and may in fact be essential to the production of the data upon which revisionist arguments rely, it is far from clear (Section 6) how anything like the ho-hum conditions on knowing could be explained away by a revisionist semantics together with the sort of pragmatic considerations adduced here. Thus, the present discussion gives us no reason to suspect the ho-hum view itself. On the contrary, that view emerges as even better-supported than it was before.

2. IMPLICIT EPISTEMIC COMMITMENTS -- ASSERTION AND THE HO-HUM VIEW

You can learn much from others' saying what they do – much more than is to be found in the words they actually utter. Gricean (1989) implicatures – where a speaker means what he says and communicates something else besides – are an instance of this. But so are Sperber and Wilson's (1986) and Carston's (1988) 'explicatures', Recanati's (1989, 1993) 'strengthenings' of sentence meanings, and Bach's (1994) 'implicatures'. – These are all ways of picking out what is communicated by an utterance, where what is communicated is an 'expansion', 'development', or 'completion' of the sentence actually uttered.² To take a well-worn example, an utterance of 'I've had breakfast' will typically communicate that the speaker has had breakfast *that day*, not the weaker proposition, simply, that he has had breakfast [at some point or another]; and "You won't die," said by a mother to her injured charge, is liable to be meant to communicate that a scraped knee isn't serious, for example, not that the child is immortal (Bach 1994).

How does this work? According to one very well-entrenched view, what enables speakers to communicate things over and above, and at times quite other than, what the words they utter mean is the fact that our conversational exchanges are governed by something like Grice's (1989) Co-operative Principle (CP). Or, better, they are governed by the mutual presumption that others try to conform to CP – hence, that they say what they do with the intention of communicating information that is, in a nutshell, maximally relevantly informative (Harnish 1976). Nor is the requirement of relevance gratuitous. For it needn't be the obvious literal falsity of what is said (cf. the mother's statement), for instance, that triggers the inference to the information the speaker intends; it could just be – as it is in the breakfast case – a lack of relevant specificity (Bach 2000, p. 265). Lastly, the two examples above show that whether a speaker conforms to CP is not to be determined by asking whether the sentences they utter are themselves maximally relevantly informative (Grice 1989, pp. 33-34). The sentence the breakfast-eater utters is too non-specific; what the mother says is literally false. But they're both being co-operative, inasmuch as they say what they do in order to get across something that is both true and maximally relevant.

In both of these examples, of course, the information made available by the relevant utterances is plausibly *intended* by the speaker to be inferred.³ But even when a bit of information isn't specifically intended by the speaker, it may be something to which he commits himself in uttering *U* with that specific force and content.⁴ So, for instance, a speaker who asserts that *p* commits himself as to *p*. Once again, that the speaker is so committed is no part of the content of the uttered sentence, '*p*'. But as the example of 'Moore's paradox' – "*p*, but I don't believe it" -- makes especially clear, in asserting that *p*, that one believes that *p* is something that is conveyed 'for free', and regardless of whatever specific communicative intentions one might have. Further, Moore's example demonstrates that our response to a given utterance can be a product of what the speaker thereby represents as being the case, whether or not the latter is reflected in the literal content of what he says. Because of these two things, an utterance such as Moore's persists in causing discomfort even in those who are quite aware that it involves no semantic inconsistency.

It's clear, of course, why asserting should involve taking on the extra-semantic commitment upon which Moore's paradox depends. For if our talk is governed by the CP, then 'saying' itself presumes one's striving to fulfil certain credal-epistemic conditions: chief among the Gricean maxims is that of Quality, 'Try to make your contribution one that is true,' along with its two more specific sub-maxims:

- i. 'Do not say what you believe to be false;' and
- ii. 'Do not say that for which you lack adequate evidence.' (Grice 1989, p. 27)

No surprise, then, that a speaker's taking on a commitment as to what he asserts is inseparable from his asserting at all.

For the sake of having a label for the phenomenon we've been describing, let's say that a speaker *conveys* that to which he commits himself -- what he represents as being the case -- in virtue of uttering '*p*' with that specific force and content (cf. Sainsbury 1984).⁵ Among the things to which the speaker so commits himself, of course, will be things (semantically) entailed by the words he utters. But, as the previous example shows, what his utterances convey can include much that's not plausibly viewed as so entailed. Using ' $\rightarrow$ ' to indicate this notion of conveying, we can represent the lesson of that example as

follows:

- (1) $S: "p"_{\text{assertion}} \rightarrow S \text{ believes that } p.$

Nor does (1) exhaust what's conveyed in asserting *per se* -- the commitments perforce taken on. For an assertion *just is* a type of presenting-as-true: the speaker's committing himself to the truth of p is part-and-parcel of his performing this speech act at all. So we can say too,

- (2) $S: "p"_{\text{assertion}} \rightarrow S \text{ believes truly that } p.$

(Remember – here, and throughout the present discussion – that ‘ $\rightarrow$ ’ is being used to mark the commitments one takes on in asserting various things. (2), e.g., does not state that, given that S asserts that p , it follows that S has a true belief that p , but only that he so represents himself.)

Further, if the speaker is presumed to be conforming to the maxim of Quality, from the fact that he asserts that p it can also be inferred, via the second sub-maxim, that he takes himself, at least implicitly, to be justified in believing that p – or, in Grice's terms, to have adequate evidence that p . Of course, just as there may be insincere assertions, one can assert that p , expressing thereby a belief that p , without in fact having any real reasons, grounds or justification for that belief. But this does not affect the point that, in asserting, one represents oneself as having such reasons (grounds, etc.). And that one does so represent oneself is confirmed by the oddity of saying, e.g., “ p , but I have no reason whatever for thinking that p .”⁶

Hence,

- (3) $S: "p"_{\text{assertion}} \rightarrow S \text{ is justified in believing that } p.$

Combining (2) and (3), we now have,

- (4) $S: "p"_{\text{assertion}} \rightarrow S \text{ has a justified true belief that } p.$

Which, on a crude, ‘justified true belief’ (JTB) account of knowing, gives us,

- (5) $S: "p"_{\text{assertion}} \rightarrow S \text{ knows that } p.$

Of course, no epistemologist regards the crude JTB theory as wholly satisfactory. Though it might take some *Theatetan* prompting on our part to get them to explicitly acknowledge it, JTB and knowledge

are close to equivalent for ordinary folks. Even so, those who've read their Gettier know that the JTB theory needs somehow to be strengthened. But then, that the speaker's belief is not only justified and true, but ungettiered, is implicit in his asserting that *p* as well. For were the speaker to become aware of the fact that, if that to which he attests is true, it is true only by accident (one standard characterization of what goes wrong in Gettier cases), he may well cease to believe; and even if not, the CP would recommend that he withdraw his assertion, as he would in effect have learned that his evidence was, in this sense, not adequate after all. It therefore seems that, in asserting that *p*, one represents oneself as satisfying the conditions on knowledge, as laid down by *the ho-hum view* – i.e., the view that knowledge is ungettiered justified true belief.

So stated, the ho-hum view is extremely uncontroversial (hence its name) – the majority by far of extant theories of knowledge can be seen as specific attempts to elaborate the ho-hum view in a satisfactory way. Of course, epistemologists disagree about just how, formally speaking, the Gettier problem should be addressed -- whether doing so requires a fourth condition (and if so, which), as opposed to the correct third condition. The ho-hum view itself is neutral on this score, as is the idea that asserting involves representing oneself as (ho-hum) knowing. Likewise, Gettier aside, epistemologists disagree as to just what sort of 'tying down' (Plato 1981, p. 86 [*Meno* 98a]) converts true belief into knowledge. Indeed, some theorists of the externalist stripe avoid the use of 'justified' altogether here, preferring to speak of knowledge as 'warranted' (Plantinga 1993), 'apt' (Sosa 1991), or 'entitled' (Dretske 2000) true belief. However, the ho-hum view itself is neutral as to the internalism-externalism debate: the reference to justified belief, as it occurs in stating the ho-hum view, is really just a marker for whatever further property must be added to convert an ungettiered true belief into knowledge. In more deliberately neutral terms, we could characterize the ho-hum view as the idea that knowledge is ungettiered true belief, together with the subject's *being in a good epistemic position* with respect to the proposition in question. But then, the same ecumenicalism would recommend that we be similarly neutral in stating what can be inferred from an assertion, given that the subject is presumed to be conforming to the second sub-maxim of Quality:

(6) $S: "p"_{\text{assertion}} \rightarrow S$ is in a good epistemic position with respect to p .

So we once again preserve the result, which many philosophers have wanted independently to endorse,⁷ that asserting essentially involves representing oneself as knowing.⁸ (In Searle's terms (1969 64-6), it involves conveying that the sincerity, preparatory and essential conditions for asserting are met.⁹)

Finally, insofar as the goal in presenting the ho-hum view is to state, if only in outline, the majority view of the semantics of 'knows,' a restriction must be placed on how the ho-hum conditions are to be elaborated. For, the same majority of epistemologists who are working within the boundaries of the ho-hum view are also non-sceptics, who hold that our epistemological theorizing should be guided by an attempt to preserve the thought that we do know many things¹⁰ -- that it preserves our intuitive anti-scepticism is standardly taken to be a *desideratum* of a satisfactory theory of knowledge. Hence the familiar thought that the belief and justification conditions in particular should not be given too strong a reading: we seldom possess anything like conclusive evidence, infallible justification or a perfectly reliable basis for those beliefs which we take to constitute knowledge. And many of those same beliefs fall short of anything like complete certainty or the absence of any possible doubt. Knowledge, then, may require (confident) belief, or the subject's being reasonably sure that p ; but their intuitive non-scepticism leads the majority of theorists to avoid making the belief condition much stronger than that.

But here too the parallel with the commitments taken on in asserting that p is preserved: for there is no reason to suppose that, in asserting that p , one represents oneself as being absolutely certain that p , or as having conclusive evidence, infallible justification or a perfectly reliable basis for that belief. On the contrary, that people regularly assert in the absence of meeting such very strong requirements, and that they are not thought for that reason to be asserting improperly, suggests that the belief and justification they represent themselves as having are not stronger than most theorists' moderate, non-sceptical readings of the belief and justification conditions on (ho-hum) knowing. At the same time, though, in those cases where one's degree of confidence or the quality of one's evidence (justification, grounds, etc.) falls clearly short of the sort of belief or justification commonly thought by such theorists to be necessary for knowledge, it is

also natural and common to qualify one's assertion ("I think that....;" "It seems to me that....;" "....., though I don't know for sure;" etc.), rather than to assert outright. And this suggests that the sort of belief and justification one represents oneself as having when one asserts are not obviously weaker than the sort required for knowledge. – That, indeed, is what gives such qualifications their point.

Of course, it could turn out that the anti-scepticism which typically accompanies the ho-hum view is untenable: it's perfectly possible, at least in principle, that even such a widely shared pretheoretic judgment as this should actually be mistaken. According to Peter Unger (1975), for example, it's not merely asserting that carries with it a commitment to knowing; one just as much represents oneself as knowing that *something* is the case in performing a whole host of illocutionary acts. And, given that he takes¹¹ himself to be in possession of a general argument for the correctness of scepticism, Unger regards the pervasiveness of our epistemic pretensions as showing that our talk is fraught with misrepresentation.

Now, while the cases that will concern us most here are a certain subclass of assertions, Unger may well be right that it's an essential feature of linguistic communication as such that speakers regularly though inexplicitly represent themselves as knowing various things. The lesson to be drawn from this here, however, is different from Unger's. For as we'll see over the next two Sections, when we bring an appreciation of the foregoing sort of reflections to a consideration of our use of sentences of the form, 'S knows [/doesn't know] that *p*,' we end up with a natural way of accommodating certain data that have shaped recent epistemological theorizing, without taking them to disclose anything about knowledge *per se* at all. More specifically, if we assume that the ho-hum view of knowledge is on the right track, we can use that semantics, together with pragmatic considerations, to predict the very phenomena which have been taken to support the rather more controversial revisionist claims mentioned above, including the data Unger cites in arguing that our intuitive anti-scepticism is in fact untenable.

3. SPEAKING OF KNOWING -- EXPLICIT AND IMPLICIT COMMITMENTS

As already indicated, it's being assumed here that what makes it possible for so much of what goes on in

linguistic communication to be inexplicit is the fact that our conversational exchanges are governed by something like Grice's CP. On its own, of course, no such general principle, just because it is so general, will suffice for interpreting a given utterance – one also needs to bring to bear other information one has. But here too, in a particular case there will be shared assumptions and mutually available facts. For instance, in a given case certain things will be mutually obvious – too obvious to require actually mentioning them; and letting the already-obvious remain inexplicit is partly constitutive of communicative rationality. (The maxim of Quantity enjoins us to make our conversational contributions informative, but not overly so, that of Relation to make them relevant.) Thus, if you and I, two typical English speakers, are walking through a tree-filled park and I remark, "That's a tree," you might find yourself at a loss as to why I am saying that. Not that what I say is controversial -- you and I both know perfectly well that the indicated object is a tree; it's just that the obviousness of what I say makes my report *prima facie* pointless. In such a case, the natural inference for you to draw is that I'm remarking on the obvious in order to get across something over and above, or perhaps quite other than, the information literally expressed by the words I utter (for instance, that the indicated tree is particularly impressive in some respect).

Another case of just this sort of thing involves the use of 'looks,' as in "The apples looks red." If you and I are looking at a basket of red apples in normal conditions (no red lights, etc.), that they look red will be mutually obvious. In fact, making it explicit – *saying*, "The apples look red" – is liable to put things in the wrong light. As Grice points out, saying, "X looks *F* to me," pragmatically implies that one doubts or wishes to deny that *X is F*, even though the sentence itself doesn't mean this. For, since there is a general principle enjoining us to make a stronger rather than a weaker statement in the absence of a reason not to do so (1961, pp. 140, 132)¹² – else, one is being less than maximally relevantly informative --, from the fact that one says merely that something looks *F* it can be reasonably inferred that one doesn't take it to *be F*. (If the speaker thought it *was* red, we suppose, that's what he would have said.) All this, without packing into the truth conditional content of the relevant sentences all of the information that utterances thereof are liable to carry, and thus without having to deny the plausible thought that it is because the apples are red that they

look red, and because they look red that I know that they are.

One final example. According to (1) above, asserting that *p* already involves representing oneself as believing that *p*. If the present discussion is on the right track, we should expect that speakers would exploit this fact in their communicative exchanges; and they do. For speakers very often use the explicit “I believe that...” in order, not to make explicit what would be obviously inferable from their asserting, simply, “*p*” (that they believe that *p*, or are so committed), but to communicate something else, something that’s not obviously indicative of the nature of belief – for example, that they aren’t absolutely certain that *p*.

In each of the foregoing examples, speaker exploit the fact that certain information would already be mutually obvious were they to assert that *p* in order to further their communicative goals. (It’s worth noting, then, that each of these examples is liable to strike us as most natural when it involves some use of stress – “*That’s a tree*,” “*X looks red*,” “*I believe that p*.” For such emphasis can serve to signal the speaker’s intention to communicate something beyond what the words they utter express (Grice 1989, p. 51).) Because in such a case intending to communicate *that* information would clearly involve a violation of CP, speakers are able to use the corresponding term(s) in order to communicate information that’s *not* plausibly regarded as part of those terms’ semantic contents. In effect, the would-be mutual obviousness of some information, *i*, conduces to speakers’ using sentences containing ‘*i*’ to communicate something beyond what the sentences themselves express.

Since, then, a commitment to (ungettiered) justified true belief, to ho-hum knowing, is already carried by a sincere assertion, and since this will be mutually known to be so, we have very good reason to suspect that speakers might use “I know...” to communicate something other than what the sentence they utter expresses, and what would already be obviously inferable from their asserting ‘*p*’ alone; that it might involve a speaker’s ‘taking a new plunge’ (Austin 1946, p. 99; cf. DeRose 2002, p. 185). Otherwise, in saying “I know that *p*,” the speaker runs the risk of needlessly making explicit what would be already mutually obvious. Much better to say that it’s a mistake on the part of theorists, an instance of ‘the

descriptive fallacy', to insist that first-person knowledge attributions can only be attempts to describe one's own (supposed) possessing of knowledge *per se* (Austin 1946, pp. 103, 78-79).

What else might a speaker be trying to communicate in claiming knowledge for himself? A simple substitution on

$$(5) \quad S: "p"_{\text{assertion}} \rightarrow S \text{ knows that } p,$$

gives us:

$$(7) \quad S: "I \text{ know that } p"_{\text{assertion}} \rightarrow S \text{ knows that } (S \text{ knows that } p). \text{ (Cf. DeRose 2002, pp. 185-186).}$$

But it's not clear that their having such second-order knowledge *per se* – their having an ungettiered justified true belief that they have an ungettiered justified true belief -- is a plausible candidate for what speakers are aiming to communicate when they claim knowledge for themselves. For one thing, it's not clear how often in our epistemic discourse we actually take an interest in others' (second-order) mental states themselves, as opposed to what we should think about some worldly matter, *p*. Further, it's plausible to suppose that when one represents oneself as justifiedly (and truly) believing that *p*, that one justifiedly (and truly) believes that one justifiedly (and truly) believes is already implied as well. For, as we've seen, that one satisfies the relevant credal-epistemic conditions is a commitment which derives from one's presumed conformity to the CP. And, as the CP governs rational communicative behaviour in general, there's no reason to suppose that it applies, and so generates such commitments, only with regard to what one says but not with regard to what one merely represents as being the case.

A more plausible suggestion as to what speakers might be aiming to communicate when they claim knowledge for themselves, and an idea that has appealed to many, is that "I know..." can function "as an emphatic variant of '[I] believe'" (Quine & Ullian 1978, p. 14), conveying one's special confidence that *p*, over and above the belief that *p* which would already be inferable from one's asserting, simply, '*p*', and which is required for one's ho-hum knowing that *p*. And when a doubt as to *p* or the speaker's entitlement to believe that *p* are 'in the air,' such an expression can serve as an attempt to put that doubt to rest

(Malcolm 1986, p. 212) and to suggest, indirectly, that p is something the audience should believe as well. On this view, as in the cases discussed above, speakers exploit the fact that certain information (here, that the speaker takes himself to know that p) would already be obviously inferable from his asserting, simply, “ p ”; and they add the “I know that...” precisely in order to get across the further information that they have no doubt as to p , that they’re *certain* (i.e., especially confident) that p , that they have not just “adequate evidence” (Grice 1989, p. 27) but “strong or conclusive evidence” that p (*ibid.*, p. 53), that “in a certain sense appropriate to the kind of statement (and the present intents and purposes), [they are] able to *prove* it” (Austin 1946, p. 85). Thus,

- (8) S : “I know that p ”_{assertion} $\rightarrow$ S is certain [i.e., especially confident] that p .¹³

However, insofar as it’s information of this sort – as against their merely ho-hum knowing – that speakers are concerned to get across in attributing knowledge to themselves, though the sentences they utter may well concern who knows what, they, speakers, are not ‘talking about’ knowledge at all.

This possibility, of course, raises serious questions about the advisability of drawing hasty conclusions about knowledge on the basis of speakers’ self-attributions thereof. But for now, note that if S is certain (especially confident) that p , it must be that he takes himself to be able to rule out any of the salient not- p possibilities that are worth taking seriously. (If he thought that, among the latter, there were some which his evidence didn’t eliminate, he wouldn’t be certain.) Thus,

- (9) S : “I know that p ”_{assertion} $\rightarrow$ among the various not- p possibilities being considered, either S can either rule them out (his/her evidence eliminates them), or they are not worth taking seriously.¹⁴

But what about *third*-person knowledge attributions of knowledge? Thus far, only first-person attributions have been discussed; it is the mutual obviousness of the fact that the speaker would already be committed to occupying a particular credal-epistemic position were he to assert simply, “ p ”, which accounts for “I know...”’s adding an element of extra confidence to what’s communicated ((8)); and – to this point, anyway – the latter is what’s been said to give rise the further conveyances which result ((9)). Whereas, it

can seem that when someone says, “They know that p ,” since the speaker is not talking about his own epistemic situation, that utterance commits him only to the subject’s possessing knowledge. If so, then at least our third-person knowledge attributions might be guided merely by the semantic content of the sentences used therein, and so might furnish epistemologists with data unsullied by any real possibility of pragmatic interference.

Appearances to the contrary notwithstanding, however, it is in fact plausible to suppose that a speaker’s willingness to attribute knowledge to another does depend on what, over and above the content of the sentence uttered, he would thereby be implicitly committing himself to. Since knowledge implies true belief, in attributing knowledge that p to another I commit myself to the subject’s belief’s being true, and – for the same general reasons used to motivate (5) -- represent *myself* as knowing that p (cf. Radford 1966, pp. 160-161).¹⁵ Thus:

(10) S : “ A knows that p ”_{assertion} $\rightarrow p$,

(11) S : “ A knows that p ”_{assertion} $\rightarrow S$ knows that p .

But in attributing knowledge to another I also – unsurprisingly -- commit myself to certain further things about the subject’s credal-epistemic position (as I see it). As Michael Williams points out, “‘knowledge’ is an honorific term that we apply to our paradigm cognitive achievements” (2001, p. 40). As Brandom (1995) puts it, when one attributes knowledge, one attributes a certain epistemic *entitlement* to the subject; and this is so whether that subject is oneself or another person. Thus, we can say,

(12) $S(A)$: “ A knows that p ”_{assertion} $\rightarrow A$ believes that p , and is entitled to that belief.

Now, insofar as we’re assuming the ho-hum view – insofar as we’re taking it that knowledge involves (ungettiered) true belief plus, in our most ecumenical terms, S ’s being in a good epistemic position with respect to p --, (12) may be seen as expressing, or as following very closely upon, a semantic relation. Even so, *given* (12), and without assuming anything else about what knowledge itself involves or requires, we have very good reason to expect that attributing knowledge can generate all sorts of *further* commitments. For consider what might be inferred, in a given situation, from my explicitly *saying* that S is

in a good epistemic position with respect to p , or that S is entitled to his belief. Suppose, for example, that there were some real question as to whether the subject could eliminate a certain not- p possibility, q , which it was important to us not obtain – whether his epistemic position with respect to p were really *that* good. Like ‘justified,’ ‘warranted’, and so forth, goodness of epistemic position is a matter of degree. So what do I mean, exactly, in saying that he’s in a good epistemic position? That his epistemic position is good enough for him to know? Perhaps; but whether he satisfies the ho-hum conditions on knowing needn’t be what addresses our concern – viz., whether he can rule out q , the possibility we are concerned be eliminated. (While it’s a possibility we’re taking seriously, q might not be something that S needs to be able to rule out in order to know: as we saw above, our intuitive anti-scepticism suggests that ho-hum knowing doesn’t require that one’s epistemic position with respect to p be maximally good, so to speak.) There is a presumption, however, that I strive to make my conversational contributions maximally relevantly informative. And so if I were to say “He’s in a good epistemic position,” I’m likely to be taken to mean the stronger (see n. 12) proposition that the subject’s epistemic position is good enough that he can eliminate q . (Either that or, contrary to what’s being assumed, that I don’t think q is worth taking seriously after all.) For, given our conversational concern, only this interpretation preserves the assumption that I’m striving to conform to CP – in particular, that I’m observing Relation. In this way, the entitlement implied by attributing knowledge can be used to motivate the third-person counterpart of (9), above. Thus,

- (13) S : “ A knows that p ”_{assertion} $\rightarrow$ among the various not- p possibilities being considered, either A can either rule them out (his/her evidence eliminates them), or they aren’t worth taking seriously.

Or, more generally,

- (14) S : “ A knows that p ”_{assertion} $\rightarrow S$ is in a good epistemic position with respect to p , given the contextually operative standards, insofar as they are appropriately in play.

Of course, since as we’ve just noted the goodness of epistemic position required for ho-hum knowing might not imply what’s on the right-hand side of (13) or (14), I might be speaking truly were I to

say, e.g., “He knows that p , but is not able to rule out q .” But this attempted cancellation is apt to be at least somewhat uncomfortable, as Grice allowed that some attempted cancellations are liable to be (1989, pp. 42-6; 1961, pp. 129, 137-8; cf. Rysiew 2001, pp. 494-498; Weiner 2006). For why would I have gone to the trouble of seeming to commit myself, via the granting of an entitlement, to the stronger proposition, only (apparently) to then try to take that back? Why didn’t I just say, for example, “He can’t rule out q ”?¹⁶

The more general point here is this: given the presumption of relevance, and given that attributing knowledge involves ascribing an epistemic entitlement, in attributing knowledge to S the speaker takes on commitments as to S ’s epistemic position (/the status of S ’s beliefs) that go beyond what the ho-hum conditions on knowing might themselves require. Nor is this restricted to the commitments described by (13) and (14). For just as I presumably wouldn’t grant the subject an entitlement if I thought there were some serious not- p possibility which he couldn’t eliminate, or if I thought he was not in a good epistemic position given the standards with which we’re operating, neither would I do so if I thought he should in fact be in doubt as to whether p , if I thought he shouldn’t be quite confident that p , if I took him not to have ‘the right to be sure’ (Ayer 1956), if I thought he had violated his ‘epistemic duties’, or if I thought he was being ‘epistemically irresponsible’. Thus, while the subject’s satisfying such conditions may well not be required for knowing, on the non-sceptical ho-hum view, that he does satisfy them is a commitment I’m liable to take on in attributing knowledge to him as well. Further, since our explanation of how these commitments get generated relies only on the fact that ascribing knowledge involves attributing an epistemic entitlement – something that’s true whether subject and attributor are two persons or one – they hold just as much for first- as for third-person attributions:

- (15) $S(/A)$: “ A knows that p ”_{assertion} $\rightarrow$ it is permissible for A not to doubt that p , for A to be certain [especially confident] that p ; A has ‘the right to be sure’ (Ayer 1956, p. 35) that p ,
- (16) $S(/A)$: “ A knows that p ”_{assertion} $\rightarrow$ A does not violate his epistemic duties in believing that p ,
- (17) $S(/A)$: “ A knows that p ”_{assertion} $\rightarrow$ it is not epistemically irresponsible of A to believe that p .

Taken together, the preceding reflections pose a serious challenge to theorists who wish to derive

certain non-ho-hum semantic conclusions on the basis of our ordinary knowledge-attributing practices. For it might well be that a willingness to attribute/deny knowledge depends just as much on one's willingness to undertake/foreswear commitments such as those described in (8)-(9) and (12)-(17) as it does on one's regarding the subject as possessing knowledge *per se*. Indeed, this seems like something we should expect if, as seems plausible, speakers are not very practiced at distinguishing among such factors in the first place, and so are apt to confuse what they merely represent as being the case (convey) in saying what they do with what they literally express.¹⁷ In short, we have good reason to think that when ordinary speakers attribute/deny knowledge, their concern may not always be to answer the question, "Does A know that *p*?": whether or not they themselves see the matter in just these terms, they may make claims couched in the language of knowing in order to communicate the stronger (n. 12), more specific information which their utterances can be expected to convey. Such information *may* bear a semantic relation to 'know(s)'. But whether it does cannot be decided on the basis of whether it is carried by the relevant attributions; for that it is so carried is predicted given the assumption merely of the ho-hum view plus a broadly Gricean view of linguistic communication. Of course, (8)-(9) and (12)-(17) raise special problems in particular for the revisionist arguments mentioned at the outset. Before turning to that, however, there are a few points which it will be useful to stress.

First, none of (8)-(9) or (12)-(17) is, in itself, controversial -- at least, the advocates of the revisionist claims being considered here don't deny that knowledge attributions (can) convey these things. What's controversial, and what's at issue here, is whether these conveyances are directly revealing of the meaning of 'know(s),' the truth conditions of the relevant sentences, and/or the nature of knowledge itself.

Second, on the present account, what's conveyed by a given utterance is indeed closely tied to semantic value of sentence uttered: for instance, if assertion didn't already presuppose the fulfilment of the very conditions which, on the ho-hum view, are required for the truth of 'I know that *p*', (7) would not hold; and if -- again, on the ho-hum view -- '*S* knows that *p*' didn't entail that *S*'s belief that *p* is justified/warranted/etc. (that *S* is in a good epistemic position with respect to *p*), uttering that sentence

wouldn't convey that *S* is entitled to that belief ((12)), and the further things inferable therefrom. Nevertheless, the only semantic facts that have been presumed in deriving the relevant conveyances are those licensed by the familiar and widely-held non-sceptical ho-hum view – viz., that *S* has a justified/warranted/etc. (ungettiered) true belief, where the belief and justification conditions are presumed not to require too strong a reading. We need suppose nothing further about what knowledge involves – hence, about the truth conditions of '*S* knows that *p*' – in order to generate these predictions as to the commitments taken on in attributing knowledge. All that's being resisted, then, is the thought that in order for a term's (/sentence's) content to play an essential role in the conveying of some information, the latter must be (semantically) included in or implied by the former.

Finally, it's not being claimed that all or any of (8)-(9) or (12)-(17) will necessarily capture what, on a given occasion, proves to be the most pertinent bit of information conveyed by a given knowledge attribution, and what the speaker intends to communicate.¹⁸ It may be, for instance, that in certain cases what is communicated will be simply that *S* stands in whatever position it is, exactly, which constitutes knowing that *p* – period. If I am teaching an epistemology class, for instance, and articulating the non-sceptical view that most epistemologists endorse, I might begin a lecture by enumerating some of the many things I take myself to know (my name, age, place of birth, etc.). If the stage-setting is right, what I communicate will be just what the relevant sentences express, no more and no less. What is being claimed, however, is that cases like this may well be the exception, and are not representative of the examples upon which revisionists rely; and, as we shall see in Section 5, when the latter examples *are* construed along the lines of that just given, it becomes harder to think that they really do support the relevant revisionist claims.

4. SOME ILLUSTRATIONS -- SCEPTICISM, INFALLIBILISM, CONTEXTUALISM

Consider the sceptical argument by Peter Unger alluded to in Section 1:

- (a) If someone *knows* something to be so, then it's all right for the person to be absolutely *certain* that it is so.

(b) It's never all right for anyone to be absolutely *certain* that anything is so.

(c) Therefore, nobody ever *knows* that anything is so. (1975, p. 95)

The argument appears valid, and balking at the second premise seems beside the point. For even if it's all right to be absolutely certain – to have the attitude that no possible evidence could ever affect one's confidence as to *p* (*ibid.*, p. 105) – about, say, simple logical truths, the majority by far of our beliefs will fail to constitute knowledge; and that's scepticism enough not to be worth wanting. So what about the second premise? According to Unger, it is licensed by none other than “the meaning of ‘know’ and...our concept of knowledge” (*ibid.*, p. 103) -- as we can see by reflecting on our everyday knowledge-attributing practices.¹⁹ For example, Unger writes that in saying to someone whom we know to be certain of something, “How can you be *certain* of that?,”

...we manage to imply that it might not be all right for him to be certain and imply, further, that this is because he might not really *know* the thing....Neither ‘know’ nor any cognate expression ever crosses our lips in the asking. We are able to imply so much, I suggest, because we all accept the idea that, at least generally, if one does know something then it is all right for one to be certain of it—but if one doesn't then it isn't. This suggests that there is some analytic connection between knowing, on the one hand, and, on the other, its being all right to be certain. (*Ibid.*, p. 98)

Moreover, Unger argues, that knowing entails its being all right to be certain follows from “the fact that knowing entails, at least, that one *is* certain” (*ibid.*). And *that*, in turn, is “made quite plain by the inconsistency expressed by sentences like ‘He really *knew* that it was raining, but wasn't absolutely *certain* that it was.’ Such a sentence,” Unger continues, “can express no truth: if he wasn't certain, then he didn't know” (*ibid.*). To the idea that knowing entails being certain, we need only add the premise that “certain” is ‘an absolute term’ – that to be less than absolutely certain is not to be certain at all (*ibid.*, pp. 114-118) – and we obtain the first premise of Unger's argument (and thence the conclusion, insofar as the second premise is true, or true enough, to imperil the bulk of our ordinary claims to know).

But we really should take a closer look Unger's data. For instance, suppose it is true that we can

imply that someone doesn't know by asking, "How can you be *certain* of that?" What explains this?

According to (12)-(17), above, attributing knowledge that *p* to another involves conveying such things as that the subject is entitled to that belief, that they can rule out any not-*p* possibilities are worth considering, that it is not 'epistemically irresponsible' of them to believe that *p*, and that they have 'the right to be sure (certain)' that *p*. However, while attributing knowledge conveys (i.a.) that one takes the subject to have a 'right to be sure', to ask, "How can you be *certain* that *p*?" is pretty clearly to imply that perhaps the subject *shouldn't* be sure that *p*. Why might such a challenge to the subject's 'right to be sure' imply that he doesn't know?

On the face of it, it might seem that it could not. Among the things which the speaker represents as being the case in saying "*p*" will, of course, be whatever is (semantically) entailed by the words he utters; and it is easy to see why, in denying something that is so entailed the speaker is able to imply that not-*p*. On the other hand, consider the most familiar examples of conversational implicatures – e.g., a reference letter writer's confining himself to remarking on the fact that the candidate for a philosophy job has beautiful handwriting (cf. Grice 1961, pp. 129-130). In so doing, the writer represents it as being the case that the candidate is not well-suited for the job. But it's clear that, were he to write instead that the candidate *is* well-suited, he would not thereby imply anything about the latter's handwriting.

However, two things need to be noted. First, from the point of view of the current discussion, it's far from clear that familiar implicatures such as the foregoing are the right model to focus on here (Rysiew 2001, p. 510, n. 32; 2005, p. 62, n. 13). For in implicatures, properly so-called, one means what one says but also something else; whereas, among the central issues before us is whether in fact knowledge attributors typically do mean just what the sentences they use express, as opposed to the various further things which they convey in using them. If and when they do read the latter information onto those sentences, they will fail to distinguish between denying what would be conveyed by attributing knowledge and denying that the subject knows.

Second, there do seem to be clear cases in which an utterance of '*p*' does or would convey that *q*,

one can cast doubt on p itself by denying q , and yet, as in the letter-writer case, p doesn't entail q . For instance, if I say, "I simply don't believe that p ," I am explicitly denying what would be conveyed by my asserting ' p ' – viz., that I believe that p (3). But in saying this I am also able to cast doubt on p itself, implying that I think it is not the case. What enables me to do this, it seems, is the fact that, here, there is an especially close tie between what would be conveyed by my asserting ' p ' and what I actually say. For whereas believing that someone is not well-suited for a certain job isn't inferable merely from the fact that one asserts that they have beautiful handing (it's only given quite specific features of the context that, given the latter, conformity to CP requires the former supposition),²⁰ that one believes that p is presupposed by one's asserting that p in the sense described in Section 2: though one may assert insincerely, given one's presumed conformity to the CP, that one does believe that p is a commitment one takes on merely in virtue of asserting that p . And, plausibly, it's because I am explicitly disavowing such a commitment -- denying something that would in this way be presupposed by my asserting ' p ' – that, in saying, "I don't believe that p ," I am able to imply something about p itself (as I see it, of course).

This too suggests an answer to the question of why, in challenging a subject's 'right to be sure', one may imply that he doesn't know. For among the reasons one might issue such a challenge is because one thinks that there is some real question, or grounds for doubt, as to p . But insofar as a speaker's challenge to the subject's right to be sure indicates the speaker's own doubt as to p , it will also thereby serve to call into question whether, by the speaker's lights, the subject knows that p ; for the subject's knowing that p requires that p be true, and the speaker's attributing knowledge to the subject presupposes that he (the speaker) *believes* that p is true (10).²¹ Especially given such a close presuppositional tie, our general tendency to read things which are merely conveyed onto the words actually uttered is liable to have us 'hearing' the speaker's challenge to the subject's knowledge in his questioning the latter's 'right to be sure'.

Next, while utterances of sentences such as, 'He really *knew* that it was raining, but wasn't absolutely *certain* that it was,' can indeed sound strange, it needs to be argued that this is due to semantic factors. And that it is not is suggested by the existence of cases in which the relevant utterances don't sound

so odd. – This is a point to which we’ll return in the next Section. But for now, suppose for example that we are fallibilists, who regard knowledge as compatible with the possibility of error, in the sense that one can know on the basis of non-entailing evidence. Now, suppose we have subjects who do know that *p*, and so do believe that *p* with whatever degree of confidence is required for knowing, but who happen at the moment to be thinking hard about those various ways in which they might be mistaken (such possibilities being, by hypothesis, compatible with their knowing). If so, they’re liable not to be absolutely certain. In such a case, we would be speaking truly in saying of the subjects, “They know, but they aren’t (absolutely) certain.” So, *pace* Unger, there appears to exist at least one scenario in which such a sentence might express a truth.

Finally, notice that while Unger is arguing for claims about the meanings of the relevant terms, nearly all of his examples involve the use of stress, which he himself regards as an extra-semantic factor – something that “does not alter the content of what is said” (1975, p. 79). Thus, for instance, Unger speaks of “the inconsistency expressed by sentences like ‘He really *knew* that it was raining, but *wasn’t* absolutely *certain* that it was’” (*ibid.*). But why the stress on ‘knew’ if the inconsistency here is supposed to be in the sentence itself? – In terms of what the argument is meant to establish, one might think, that stress has no business being there.²² On the other hand, recall the naturalness of speakers’ using stress in the examples of Section 3. It was suggested above that in such cases speakers exploit the fact that certain information would already be obvious, and use the stressed phrase in order to signal their intention to communicate information that’s *not* plausibly regarded as part of the relevant expressions’ semantic content. Thus, to take an example much closer to Unger’s own, consider someone saying to her spouse, “You *know* we can’t afford a new car.” Here, the stressed ‘know’ is naturally read as conveying the very sort of certainty (or the entitlement thereto) to which (8) and (15) refer; it has the force of something like, ‘surely you don’t doubt that....’ -- Though of course that’s not what the uttered sentence means: knowing entails truth; not having any doubts does not. And while it’s being assumed that knowing that *p* requires believing, and whatever freedom from doubt that that requires, whether it requires regarding *p* as obvious in the way or to the extent

that the speaker in this example represents it as being, much less anything like the absolute confidence that Unger supposes, is just what's at issue, and so not something that should be presumed. Plausibly, though, it's just this use of 'know(s)' that Unger is relying upon in generating the appearance of inconsistency: to an unbiased ear, the total effect of his imagined utterance is something like, "They were *sure* that *p*, but not *certain* that *p*." Here, the inconsistency is very close to the surface, but nothing follows about what knowing entails.

But what about fallibilism? On the face of it, fallibilism is inseparable from the intuitive anti-scepticism we were just defending, and which typically accompanies the ho-hum view. If knowledge requires the impossibility of error (evidence that logically entails what's believed; 'certainty' in this sense), we have much, much less knowledge than we commonly suppose. So much the more troubling that there can seem to be something wrong-headed, even incoherent, about the fallibilist view itself. While this has been noted by a number of writers, David Lewis provides the most recent and eloquent statement of the worry:

If you claim that *S* knows that *P*, and yet you grant that *S* cannot eliminate a certain possibility in which not-*P*, it certainly seems as if you have granted that *S* does not after all know that *P*. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory....If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong? (1996, pp. 549-550)

Coming from our discussion of Unger, it's clear enough how we might accommodate Lewis' data. For it may be that the reason overt, explicit fallibilism²³ "just *sounds* contradictory" is that the relevant utterances involve the speaker's *conveying* inconsistent information. In the first-person case -- e.g., "I know that *p*, but of course I might be mistaken" -- the speaker puts it across both that she is certain that *p* ((8)), is entitled to her belief ((12)), and that she (/her evidence) can eliminate any salient not-*p* alternatives that are worth taking seriously ((9)/(13)), *and* that there is some uneliminated possibility of error which (because he

has mentioned it) is both salient and something the speaker thinks is worth taking seriously.²⁴ In the third-person case – “*S* knows that *p*, but *S* cannot eliminate a certain possibility in which not-*p*” – one is naturally understood as wishing to convey both (with “*S* knows....”) that *S* is entitled to his belief that *p* ((12)), that she has ‘the right to be sure/certain’ that *p* ((15)), that she (/her evidence) can eliminate the salient not-*p* possibilities ((13)), etc., *and* (with the “but...”-clause) that there’s some real chance that not-*p* and hence that the contrary of each of these is in fact the case. So, whether or not fallibilism is by its nature incoherent, we should expect that “overt, explicit” fallibilism will sound “wrong” (*ibid.*, p. 550). For all the data show, then, it could be that what is “elusive” is not knowledge, but the conversational conditions under which an explicitly fallibilistic knowledge attribution will not convey contradictory information.

A pragmatic response to Lewis’ problem (in its first-person version) along lines similar to the foregoing has recently been suggested by Michael Williams (2001, p. 54). Likewise, the present proposal is very much in the spirit of Sosa’s proposed means of accounting for the (alleged) oddity of utterances such as, “They know that *p*, but they need to investigate further.” Sosa suggests that the oddity here stems from the fact that “knows” (typically) conversationally implies that there is no need for further investigation.²⁵ On the present account, this in turn is to be explained by the fact that that attributing knowledge conveys an entitlement to the subject ((12)). Given CP, this further conveys that the subject can rule out any counter-possibilities worth considering ((13)), that he has the right to be sure ((15)), and so forth – commitments one would not want to take on if one thought that further investigation was in fact needed. However, according to Stewart Cohen, such pragmatics-based proposals don’t stand up to close scrutiny. For, as Grice notes, pragmatic implications are cancellable, whereas “it sounds inconsistent to say, [for example,] ‘We know, but we need to investigate further’”. This suggests that the implication is semantic” (Cohen 1999, p. 60).

But as we’ve already noted (Section 3), and as we saw in connection with Moore’s paradox (Section 2), we should in fact expect that certain attempted cancellations are going to be uncomfortable. For first, we have no reason to think that speakers will be especially adept at distinguishing between what they

mean in uttering certain sentences and what the sentences they utter mean. Second, the tendency to read the former sort of information onto the latter will only be reinforced when information that is (only) pragmatically conveyed is *standardly* carried by one's saying certain things. For there, what is conveyed is especially likely to "supplant" (Neale 1990, p. 77) what is literally expressed, with the result that cancellations will sound a lot like semantic confusions. And that the conveyances described above *do* standardly hold is suggested by the fact that our account of why they hold appeals merely to a rather sparse semantics for 'know(s)' and the mutual presumption that speakers strive to conform to CP – very little in the way of specific knowledge of the context(s) in which the relevant sentences are uttered seems to be required.²⁶

Our response to Cohen's objection sets the stage for a pragmatic treatment of the data that is said to support the contextualist view which he, and a number of other epistemologists, have recently endorsed. Contextualists themselves are generally prepared to allow that contextualism is something that one needs to be argued into: it takes *work* to come to think, for example, that there can be situations in which we have two subjects, exactly alike psychologically, possessing the very same history and the same evidence with regard to *p*, etc., only one of whom is said truly to know that *p*, since they are being evaluated from within different contexts: we seem, if anything, to be 'intuitive invariantists'. (Hence contextualism's not being among the tenets of ho-hum epistemology, but a revisionary view.) And yet, contextualists argue that our own knowledge-attributing practices provide very good evidence for the view that "the truth conditions of sentences of the form 'S knows that *p*' or 'S does not know that *p*' vary in certain ways according to the context in which they are uttered" (DeRose 1992, p. 914). Thus, Keith DeRose writes that "[t]he best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, nonphilosophical talk: What ordinary speakers will count as 'knowledge' in some non-philosophical contexts they will deny is such in others" (2005, p. 172, forthcoming). Likewise, Cohen claims that examples such as the following "strongly [suggest] that ascriptions of knowledge are context-sensitive" (1999, p. 59):²⁷

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and respond, “Yes I know – it does stop in Chicago.” It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John agree that Smith doesn’t really *know* that the plane will stop in Chicago. They decide to check with the airline agent. (*Ibid.*, p. 58)²⁸

While there are important differences between their views, contextualists such as Cohen and DeRose agree that, in such cases, the practical importance of the subjects’ getting it right tends to raise the standards for the truth of a sentence of the form ‘*S* knows that *p*’. The result is that the knowledge denial (Mary and John’s claim) in the high-stakes situation may be true, without affecting the truth of the low stakes claim to know (by Smith himself).²⁹ What reason is there, though, for adopting this way of resolving the apparent incompatibility between the two knowledge utterances? According to Cohen, it is important that the low-stakes claim (Smith’s) be true, since only that preserves our intuitive anti-scepticism – the thought that we do know many things: if you can’t know on the basis of ordinary, non-entailing evidence such as what’s printed in the flight itinerary (which neither we nor Smith have any reason to think is erroneous), then we’ll have to deny very many of our ordinary claims to know (*ibid.*, p. 59). On the other hand, if John and Mary’s stricter standard is *too* demanding – if their denial of knowledge to Smith is false – then it’s “hard to see how [they] should describe their situation”:

Certainly they are being prudent in refusing to rely on the itinerary. They have a very important meeting in Chicago. Yet if Smith knows on the basis of the itinerary that the flight stops in Chicago, what *should* they have said? ‘Okay, Smith knows that the flight stops in Chicago, but still, we need to check further.’ To my ear, it is hard to make sense of that claim. (*Ibid.*, pp. 58-9)

We have just seen, however, that the oddity of the latter claim may be accounted for without resorting

to a revision of our ordinary, ‘invariantist’ thinking about knowledge. More generally, the materials of Section 3 give us a way of explaining the apparent context-dependence of knowledge sentences’ truth conditions in terms of the context-variability of the specific commitments taken on in attributing knowledge - i.e., the information on the right-hand side of any of (8)-(9) and (12)-(17) --, and/or the perceived truth-value thereof. For whatever else its semantic content, ‘*S* knows that *p*’ implies that *S* is in a good epistemic position with respect to *p*. But what people are willing to count as the subject’s being in a good epistemic position is a variable matter – in terms of (14), the contextually operative standards can shift. (If the stakes are high, an epistemic position that is, in fact, good enough for the subject to know might not be good enough for them. For Mary and John’s purposes, e.g., it matters not whether Smith satisfies the conditions laid down by the correct epistemological theory.) Because of this, the specific commitments one takes on in attributing knowledge will be variable as well. (Were Mary to grant Smith knowledge, she’d be putting it across that his epistemic position is good enough that she and John needn’t worry about whether the plane stops in Chicago and that the possibility of a misprint in the flight schedule isn’t worth taking seriously, etc. ((13)-(14)), both of which may well be false. But it’s not that Mary and John doubt that Smith believes that the flight stops in Chicago, or that they think it in fact doesn’t; nor do they doubt that consulting an itinerary is generally sufficient to put one into a good epistemic position with respect to such matters. So their denial of knowledge to Smith is liable to convey, and to be understood as intended to convey, precisely that Smith’s epistemic position isn’t good enough to satisfy the unusually high epistemic standards that are for them in play.) Add to this the fact that our untutored intuitions about the truth conditions of various sentences are generally insensitive to the semantic/pragmatic distinction -- we often read aspects of the total communicated message onto the sentence uttered in delivering it³⁰ -- and we get the result that what ordinary speakers will count as knowledge in some situations they will deny is such in others.³¹

5. FURTHER PROBLEMS FOR THE REVISIONIST ARGUMENTS -- THE HO-HUM VIEW RATIFIED

We've just seen how, given some rather widely-accepted assumptions – that assertion involves representing oneself as (ho-hum) knowing, that our conversations are governed by the presumption that speakers strive to conform to the CP, and that our ordinary 'semantic intuitions' are the product of both semantic and pragmatic factors – we are able to predict the data that have inspired proposed revisions to the ho-hum view. Of course, the discussion of the previous Section is hardly conclusive – the suggested explanation of the relevant data may prove not to be the best explanation thereof. But the possibility of such an alternative explanation is not the only problem the revisionist arguments face. For if we suppose, against the arguments of the previous Section, that pragmatic factors are not importantly responsible for producing the allegedly pro-revisionist data, those data in fact become much harder to predict.

To begin to see this, recall that none of the conveyances described in Section 3 are themselves controversial. As noted previously (Section 2), Unger for example is on record as endorsing

$$(5) \quad S: "p"_{\text{assertion}} \rightarrow S \text{ knows that } p.$$

And he would regard both of the following as holding as well,

$$(8) \quad S: "I \text{ know that } p"_{\text{assertion}} \rightarrow S \text{ is certain [i.e., especially confident] that } p.$$

$$(15) \quad S(A): "A \text{ knows that } p"_{\text{assertion}} \rightarrow \text{it is permissible for } A \text{ not to doubt that } p, \text{ for } A \text{ to be certain [especially confident] that } p; A \text{ has 'the right to be sure' that } p.$$

For, as we've seen, Unger holds that both (8) and (15) are also semantic entailments: as he sees it, the reason S represents such things as being the case in uttering the relevant sentences is because of what those sentences *say*.

But this has some puzzling consequences. For example, given Unger's preferred conception of knowledge, (5) licenses,

$$(5_U) \quad S: "p"_{\text{assertion}} \rightarrow S \text{ is certain that } p.$$

However, if asserting that p already involves representing oneself as being certain that p , what illocutionary

function is served by asserting *that one knows* that p ? On the account of Section 3, the use of ‘I know’ can serve precisely to convey that the speaker is certain (in the sense of especially confident) that p , something that’s not inferable from one’s asserting, simply, ‘ p ’ (though that one is committed to satisfying the *ho-hum* belief condition on knowing is so inferable). On Unger’s view, though, that one is certain would already be obviously inferable from the assertion, “ p ”, in virtue of the sentence’s truth conditional content. So why would speakers go to the bother of actually claiming knowledge for themselves? Insofar as they are communicatively rational, one might think, they would do so in order to communicate *some* further information. What could that be? An Ungerian version of (7),

$$(7) \quad S: \text{“I know that } p\text{”}_{\text{assertion}} \rightarrow S \text{ knows that } (S \text{ knows that } p),$$

would be,

$$(7_U) \quad S: \text{“I know that } p\text{”}_{\text{assertion}} \rightarrow S \text{ is certain that } (S \text{ is certain that } p).$$

But whereas (7) clearly involves the conveying of information not mentioned in (5), it’s not clear that (7_U) adds anything that’s not in (5_U). For ‘certainty’ here is a wholly psychological matter; and while there are some attitudes (a feeling of vague unease, for instance) such that one can plausibly fail to realize that one has them, it’s usually thought to be in the nature of (psychological) certainty that one *knows* that one is certain.

What this suggests is that, given Unger’s revisionist view of knowledge, “I know...” wouldn’t be adding anything to what would be conveyed by the bare assertion, “ p ”. And so too for third-person attributions, presumably. For if we took “He knows...” as functioning to communicate anything other than what’s expressed by the uttered words themselves (by Unger’s lights, that, *i.a.*, the subject is certain), we could no longer infer a semantic inconsistency from the oddity of utterances like, “He really *knew* that it was raining, but wasn’t absolutely *certain* that it was” (1975, p. 98). Thus, it must be that for Unger both “I know...” and “He knows...” are always used strictly literally: the speaker means exactly what the words he utters means, nothing more and nothing less; and our intuitions about the oddity or truth of certain

utterances must concern only the contents the sentences uttered as well.

Similar considerations apply to stress. If the oddity of certain utterances (e.g., “He *knew*, but wasn’t *certain*”) is supposed to reflect a semantic defect, the suggestion of the previous Section -- that stress might be functioning in Unger’s examples to signal the speaker’s intention to communicate information that goes beyond the truth-conditional content of the sentence he utters -- will have to be rejected. And, in fact, Unger sees stress as playing a rather different role in his argument: far from inviting the intrusion of information not present in or entailed by the sentences being considered, “the Principle of Emphasis,” like other aids to ‘getting proper focus,’ helps to ensure that we theorists’ attention is being directed precisely where it should be – at “the actual meaning of our words”, such that we’re more likely to “feel any inconsistencies those words may express” (1975, pp. 80-82).

The problem, however, is that it is no longer clear that Unger can predict the data that are supposed to support his revisionist claims. As noted above (Section 3), it’s no part of the present argument that knowledge attributions can never serve to communicate just what the sentences uttered themselves express. Likewise, there is no reason to deny that stress does sometimes serve merely as a way of focussing attention on what it is we’re saying. (If, e.g., I think you’re failing to appreciate the seriousness of a mutual friend’s actions, I might say, “He *stole* it.”) However, if the relevant sentences are themselves supposed to be problematic – if ‘He *knew*, but wasn’t *certain*’ is, in this respect, like ‘He was a *bachelor*, but not *unmarried*’ – the seeming inconsistency should survive the removal of the stress, as it does in the bachelor case (provided, of course, that we’re carefully attending to what is being said). To some, though, including myself, (i) “He knew, but wasn’t certain” is far less (in fact, not at all) odd-sounding; (ii) even more obviously all right-sounding is “He had an ungettiered justified true belief, but wasn’t certain”: and, in the previous Section, we considered (iii) a more fleshed-out example of subjects of whom we might consistently and truly say similar things.

Of course, Unger (and others) might not share that reaction to even the unstressed (i); and he (/they)

are liable to complain that (ii) and (iii) presume the non-sceptical ho-hum view, the viability of which is, in effect, just what Unger's calling into question. Further, and relatedly, the sceptical revisionist might press that it's not just the apparent oddity of utterances such as, "He knows, but isn't certain," which motivates his view, but the apparent falsity of the relevant modal claims – e.g., that *S* can know, even if *S* isn't certain. – The latter still sounds bad, he might say; whereas the modal counterpart of 'Moore's paradox' – "It can be the case that *p*, even if I don't believe it" – sounds perfectly all right. And *that*, it might be said, confirms that while Moore's paradox really is only pragmatic, the claims being considered here are semantically confused.

But we can sharpen the present point. For if the task before us is to focus our attention on the meanings of the relevant linguistic items, such that we can be confident that it's those meanings which are responsible for any inconsistencies we're sensing, we ought to be able, not only to abandon the use of stress, but to substitute into the relevant sentences the proposed cashing-out of the key term(s) in question. In fact, by in this way laying bear the meanings of the relevant terms, we should be *even better* able to spot any semantic inconsistencies being expressed. (As we would be with 'He was a bachelor, but not unmarried': as obviously problematic as this sentence is, it's even worse when rendered as, 'He was an unmarried male, but not unmarried'.) And here it's very important not to stray from just what it is that Unger takes certainty, in the sense required for knowing, to involve. For by '(is) certain', Unger doesn't mean merely '(confidently) believes', or 'doesn't in fact doubt'. – Either of these would be tantamount to the non-sceptical, ho-hum view's psychological condition on knowing. So if such were what Unger meant by '(is) certain,' the ho-hum theorist could agree that 'He knew, but wasn't certain,' or its modal counterpart, is indeed inconsistent; but neither of these attitudes is obviously always impermissible, as the second premise of Unger's argument has it. Nor does Unger mean even the especially confident believing to which (8) and (15) refer. -- If this were the sort of 'certainty' Unger meant we could, as indicated in the previous Section, anticipate the discomfort Unger notes, but without having to regard it as indicating any semantic defect in the relevant sentences.

As noted above, though, Unger takes ‘certainty’ to be something else again. He regards ‘certainty’ as an absolute term, and the sort of certainty required for knowledge to be the most extreme dogmatism: if I’m certain that *p*, then I’m of the attitude that I will reject as misleading, without any further thought or consideration, anything which seems to suggest that not-*p* (1975, p. 31); certainty is the attitude that “no matter what new experience I may have, I will be no less certain but that *p*” (*ibid.*, p. 105). It is the unreasonableness of *this* attitude that’s required for the second premise of Unger’s argument, and the absence thereof that’s supposed to be incompatible with knowing. But if it’s our awareness of the latter fact that’s responsible for whatever discomfort attends utterances like, “He knew, but wasn’t certain,” the sentence, ‘He knew that *p*, but was not of the attitude that absolutely nothing could ever make him any less sure that *p*,’ should strike us as at least equally bad, as should its modal counterpart, ‘It can be the case that *S* knows, even if *S* is not of the attitude that absolutely nothing could ever make him any less sure that *p*.’ But I take it that even those who find “He knew, but wasn’t certain,” odd-sounding will find the latter sentences less obviously so, if indeed they detect in them any kind of inconsistency at all. The upshot is that it’s not clear how, given his revisionist, non-ho-hum view of knowledge, Unger can accurately predict the very data that are supposed to provide evidence for that view. Whereas, with the help of (8) and (15) in particular, the ho-hum theorist can predict that the utterances upon which Unger’s argument relies might produce some discomfort; but he does so by presuming the unrevised ho-hum view. In this way, far from undermining the latter, the data in question actually support it.

A similar problem arises for contextualists who rely upon examples such as the airport case, discussed above. For such examples to provide support for contextualism, it is essential that speakers’ attributions/denials of knowledge, and our own assessment of the ‘correctness’ of what each says (Smith, in his ‘low standards’ setting; Mary and John, with their higher standards), are guided solely by the semantic facts: any suggestion that something else is responsible for these things immediately reduces the pressure to preserve the appearance that what each of the speakers says is ‘correct’ by adopting a contextualist semantics.

Let's assume, then, that the information that each of speakers in the airport example is concerned to communicate, and to which our intuitions about their utterances are responsive, is nothing more or less than that which is, according to the contextualist, what the sentence he/she utters semantically expresses. While they disagree as to how strength of epistemic position is best understood, contextualists tend to agree that a token of '*S* knows that *p*' is true just in case the subject has a true belief and is in a strong epistemic position, with standards for how strong the subject's epistemic position must be in order for the tokened sentence to be true being contextually variable (Conee 2005, p. 51, Brown 2006, p. 407). Using subscripts to denote this context-relativity, we can bring out the (alleged) truth-conditional content of the two uttered sentences in the airport case as follows:

(Smith) '*S* has a true belief and is in a strong_{C1} epistemic position,' and

(Mary) 'It is not the case that (*S* has a true belief and is in a strong_{C2} epistemic position)' (see Bach 2005, Section I).

However, while each of these sentences does indeed seem true given the details of the example, it's not clear how this provides any support for contextualism. For the two claims seem perfectly compatible, both with each other, and with a non-contextualist ho-hum view of the truth conditions of knowledge sentences. Whereas, once again, the allegedly troubling (i.e., pro-revisionist) data – viz., the apparent 'correctness' of what each speaker says in the original example -- is predicted by the unrevised ho-hum view (Section 4). So, as with Unger, if we take the pragmatics seriously, there's no reason to adopt the revisionist claim; but on the assumption that the pragmatics is really irrelevant, the revisionist has trouble sustaining the idea that data in question really do support his view. In this way, the existence of such allegedly revisionist-friendly data is once again actually evidence for the non-revisionist, ho-hum view.

Of course, contextualists can and do hold that we have a tendency to "conflate contexts" (Cohen 2005, p. 69) – that is why we can fail to see that certain low-standards knowledge attributions are, in fact, compatible with the corresponding high-standards denials (see n. 29). Again, though, this response comes at

the expense of weakening the evidence for contextualism: to the extent that our intuitions as to the truth-conditional contents of various tokened knowledge sentences are less than fully reliable, it's less clear that the best way to explain those intuitions is through the device of a contextualist semantics and the supposition that a grasp thereof is what guides speakers in their knowledge-attributing behaviour, and what's responsible for our assessments of what they say.³²

Finally, as for the Lewisian argument discussed in the previous Section, if the trouble is supposed to lie simply in the fallibilist idea itself, then it shouldn't matter what terms, exactly, are used in making that idea explicit. Thus, one should expect that an utterance of, 'He has an ungettiered true belief, based on extremely good, non-entailing evidence,' for example, would clash in some obvious way with 'He knows' – that it would do so to just the same extent, and for just the same reason, that the sort utterances which Lewis himself features ("S knows that *p*, but he cannot eliminate the possibility that not-*p*," say) are odd-sounding. But, in fact, such a clash is rather hard to detect. (That, indeed, is why it has been possible for so many to take fallibilism seriously and to suppose that 'S has an ungettiered true belief, based on extremely good, non-entailing evidence' might *imply* 'S knows'.) That there is such an asymmetry, however, poses no special problem for the ho-hum theorist, who can advert to such extra-semantic factors as the specific manner in which the fallibilist idea is expressed, and what is and is not plausibly conveyed thereby, in explaining why Lewis' utterances are odd-sounding while the clash predicted here does not obviously occur.

6. "TU QUOQUE"? ON THE STATUS OF THE HO-HUM VIEW

Equipped with the widely-held ho-hum view and an equally non-radical account of the influence of non-semantic factors on our talk, we can accommodate the data that give rise to the thought that the former stands in need of some rather non-ho-hum revision. Indeed, we've just seen how the ho-hum view may in a certain sense be necessary to the production of those seemingly revisionist-friendly data. But that, of course,

isn't why so many epistemologists have gravitated towards the ho-hum view -- the present argument may extend the explanatory power of this widely shared conception of knowledge; but it has such explanatory power quite independently of anything said here. The thinking here is available in any introductory epistemology textbook, and shouldn't need to be rehearsed in any detail. Briefly, though: just as not giving the belief and justification conditions too strong a reading allows us to preserve our intuitive non-scepticism, the ho-hum conditions generally enable us, not only to explain our tendency to deny that the subject knows when it turns out that his belief is false, or merely accidentally true,³³ but also to account for such things as the stability of epistemized belief and the role of knowledge in explaining successful action (e.g., Plato *op cit.*; Williamson 2000, pp. 60-64, 78-9; Weiner 2005, Section 4; Kornblith 2002, Chapter 2).

Of course, the ho-hum view confronts some *prima facie* contrary data. But for such familiar challenges, there are equally familiar answers. Thus, for example: Someone might say, "I don't believe it, I know it;" but such a claim is very easily understood to involve, even to require, an implicit 'merely'. People, especially non-philosophers, sometimes speak of knowledge where it's clear that the belief in question is false; but the ho-hum theorist can say that "the ascription of knowledge in such cases reflects the point of view of the people in question, from which the proposition *seems* true" (Bonjour 2002, p. 33). And, as Kent Bach observes, "most of the time, outside of epistemology, when we consider whether somebody knows something, we are mainly interested in whether the person has the information, not in whether the person's belief rises to the level of knowledge" (2005, pp. 62-63); so it should not surprise us if we sometimes attribute knowledge without looking into whether the relevant belief is justified, and ungettiered.

Still, one might wonder whether the present attempt to defend the ho-hum view, and to extend its explanatory power, doesn't actually undermine the ho-hum view itself. For while the present argument doesn't provide anything like a basis for rejecting our epistemological intuitions wholesale, it does suggest the general moral that epistemologists have tended to seriously underestimate the influence of pragmatic factors on our knowledge-attributing behaviour and our intuitive responses to various particular attributions/denials of knowledge. And just as we have (it's been argued) uncovered specific reason to be

suspicious of the data cited in support of the revisionist views discussed above, it might appear that we've also uncovered, if only inadvertently, specific reason to think twice about the ho-hum view itself. For, according to the discussion of Section 2, in asserting that p one already represents oneself as satisfying the very conditions standardly taken (by the ho-hum view) to be conditions on knowing ((5)); likewise, in attributing knowledge to another, the speaker represents himself as knowing that p ((11)). One might think, therefore, that this furnishes the non-ho-hum theorist with a convenient way of explaining away one or more of those conditions in pragmatic terms: *whatever* the truth-conditional content of ' S knows that p ,' it might be said, thanks to the conditions on sincere assertion we can expect attributions of knowledge to convey that S has an ungettiered justified true belief.

But this simply won't work. At most, the preceding line of thought would explain why those who assert, or who attribute knowledge to another, represent themselves as ho-hum knowers. It would not explain the appearance that those to whom knowledge is attributed must, in virtue of *that* fact, satisfy the ho-hum conditions. For when S asserts that p he represents himself as satisfying the ho-hum conditions *with respect to* the propositional content of what he asserts; whereas, what needs explaining if the ho-hum view is to be debunked is why we would think that the subject's satisfying the ho-hum conditions is *part of the content* of what he asserts *when he claims knowledge for himself*. Similarly, when one attributes knowledge that p to another, one represents oneself as ho-hum knowing that p ; but that does nothing to explain why we'd think that the subject of that attribution must himself satisfy the ho-hum conditions if the attribution is to be true.³⁴

And there are other problems too. For in addition to needing some way of explaining the whole range of considerations typically taken to support the ho-hum view, the proponent of such an alternative, non-ho-hum conception of knowledge would owe us some positive evidence for that conception and, relatedly, some specific reason for thinking that one or more of the ho-hum conditions on knowing should actually be rejected. Also, he'd need to account for the various further things which, according to Section 3, attributing knowledge conveys: on the present account, those conveyances are derived by presuming the ho-

hum view in particular; it is an open question whether anything like them – hence, any kind of response to the revisionist views we’ve been considering – could be arrived at given some non-ho-hum starting point.

Of course, another, more promising attempt to debunk one or more of the ho-hum conditions on knowing may be waiting in the wings; and nothing like an exhaustive defence of the ho-hum view has been attempted here. But the preceding should dispel any thought that the present argument is obviously self-undermining. Also, it serves to remind us of the range of data that would need to be explained by any non-ho-hum theory and, by comparison, of the rather impressive explanatory power of the ho-hum view itself. The latter may be mundane, but there are very good theoretical reasons for why it is so widely held.

7. CONCLUSION

Were the question put to them directly, epistemologists would pretty much universally reject the thought that semantic facts alone guide us in our knowledge-attributing behavior. Even so, we’ve just seen how some influential challenges to our ordinary ways of thinking about knowledge may require precisely this idea. The various revisionist positions treated here might, for all that’s been said, be true. But the evidence that has figured most prominently in arguing for them are certain features of our everyday knowledge-attributing practices; and, if the present arguments are correct, that evidence has now been neutralized. Meanwhile, the ho-hum view seems if anything to be in even better shape than it was before. Not only are there strong and familiar theoretical considerations lying behind it, it is untouched by the present reflections, and in fact may be essential to the very phenomena that have been thought to pose one or another kind of threat to it.³⁵

NOTES

¹Nothing much hangs on the exact details of how we draw the semantics-pragmatics distinction. But here it is assumed to mark a difference between information encoded in linguistic types and information arising out of

utterances and those facts/principles which are responsible for its so arising (Bach 1999).

²There are of course disagreements among these theorists – not least, concerning the proper conception of ‘what is said’. But they agree that what’s communicated by a given utterance is often other than the literal content of the uttered words.

³And intended to be inferred partly on the basis of the hearer’s recognition that it is so intended. That is, these examples concern information’s being made available by an utterance in virtue of being *communicated* (Bach and Harnish 1979).

⁴The latter qualification, which will henceforth be taken for granted, distinguishes conveyed information from that which is *merely* “made available” (Neale 1990, p. 74) – i.e., such information as can be reasonably inferred from a given utterance, regardless of its specific force and content. Thus, e.g., though my uttering any number of other things would do so as well, “I like ice cream” makes available to any competent hearer the information that my vocal cords are intact, that I speak English, etc.

⁵Different theorists have suggested different technical labels for what’s here being called ‘conveyed’ information. There are, e.g., the notions of *representing oneself* as having certain attitudes (Unger 1975), of *discursive commitment* (Brandom 1994, 1995), of *R’ing* (*taking responsibility for* certain things obtaining) (Alston 2000), and of *[pragmatic] presupposition* (Stalnaker 1974). The present discussion neither requires adopting any of these authors’ developed views nor presumes the priority of any of them over the others. What’s essential for present purposes is simply the idea that certain things will be obviously inferable from a speaker’s uttering ‘*p*’ (in context) with that specific force and content -- these are the things to which the speaker perforce commits himself, or which he ‘conveys’.

⁶There are, of course, related utterances which don’t sound odd, such as “*p*, but that’s just my opinion.” The naturalness of saying such things, however, would be a problem for the idea under consideration only if it were most plausible to take the speaker’s saying “but that’s just my opinion” quite literally – i.e., as indicating that the speaker really does see himself as having no reason at all for supposing that *p*. It’s at least as

plausible, though, to see “but that’s just my opinion” as a familiar device for forswearing any overly-strong epistemic commitments, perhaps because the topic is known to be one on which opinions may reasonably differ, or because one doesn’t wish to be held responsible if it turns out that not-*p*. (Similar considerations apply to “*p*, but I don’t *know* [for sure] that *p*,” and the fact that it is more natural and less odd-sounding than the straight “*p*, but I don’t know that *p*.”) Were we to discover in such a case that the speaker really did take himself not to have any grounds, justification, or evidence for supposing that *p* – that, as he saw it, he merely believed, without reason, that *p* -- we would be within our rights to criticize him for asserting that *p* and/or to regard that assertion as somehow defective. (Compare Hawthorne 2004, p. 23, n. 57, and Unger 1975, pp. 254ff.)

⁷The idea can be found, in various forms, in, e.g., Austin 1946, p. 77; Moore 1993, p. 277; Williams 2001, p. 149; Recanati 1986, p. 219; Alston 2000, esp. p. 63; Brandom 1994, pp. 203-204; Unger 1975, pp. 252ff; and DeRose 2002. (See too the further references given in the latter’s note 20.)

⁸Timothy Williamson (2000, Chapter 11) defends at length ‘the knowledge account of assertion,’ whereby one must assert that *p* only if one knows that *p*. There may be good reason to resist the knowledge account of assertion (see, e.g., Weiner 2005), but the present discussion does not require it. For one can accept the idea that in asserting that *p* one represents oneself as knowing that *p*, while regarding the knowledge rule as too strong: one could take the relevant rule to be ‘Do not assert what you take yourself not to know,’ for example, or even (as Weiner argues) ‘the truth rule’.

⁹Thanks to Kent Bach for this point.

¹⁰This idea is what Feldman (2003, p. 3), e.g., describes as “the first thesis” within ‘*the Standard View*’.

¹¹Unger, of course, has since abandoned scepticism in favor of ‘philosophical relativity’ (1984); more recently still (1986), he has moved towards contextualism.

¹²Grice is very clear about not wanting to restrict such ‘strength’ to the logical strength of the uttered sentence. (He allows that while neither of the propositions expressed entails the other, “It looks red to me” seems clearly weaker than “It is red” (1961, p. 140).) A plausible, more general notion of strength, and the

one being used here, is that of the comparative informativeness of an utterance, with a sentence's or proposition's 'strength' being a matter of how informative the corresponding utterance is liable to be.

¹³Besides its independent plausibility, the idea that claims to knowledge convey the speaker's comparative certainty is supported by the observation that, when speakers do actually say what's on the right-hand side of (7) – e.g., “I don't just know, I *know* that I know” – they are very naturally understood as intending to communicate just such confidence. There may, of course, be good reason for theorists not to equate certainty with knowing that one knows (see, e.g., Feldman 2003, pp. 126-127); but the present point concerns an apparent illocutionary function of claiming such second-order knowledge, and so a possible communicative point to speakers' representing themselves as having such, not what such second-order knowledge does or does not entail, or whether ordinary speakers have any developed views on the matter.

¹⁴The need for the second disjunct here, and in (13) below, is suggested by Brown's (2006) discussion of the roles of salience and practical importance in the pragmatics of knowledge attributions. Brown argues that the salience of a certain not-*p* alternative needn't directly affect a speaker's willingness to attribute knowledge; whether it does so will depend on the degree of practical importance of what's at issue. (Compare Hawthorne's (2004, pp. 62-66) discussion of salience, attention, and taking seriously.) Also, even if it is important to those involved to get things right, it may be that a given counter-possibility, though salient, will be seen as simply too 'unrealistic' to be worth taking seriously (Bach 2005, Section VI). For both of these reasons, and contrary to what I've previously (2001) supposed, one cannot immediately infer, from the fact that knowledge is being attributed, that the speaker takes the subject to be able to rule out all of the various not-*p* possibilities that are 'in the air'.

¹⁵Though neither discusses it in terms of what one represents as being the case in attributing knowledge that *p* to another, that the latter presupposes that the speaker himself knows that *p* has been observed by Bach (2005, Section V) and Hawthorne (2004, p. 160).

¹⁶Compare Grice on an analogous, non-epistemic case: “...the reason why it would be peculiar to say ‘She was poor but she was honest, though I do not mean to imply that there is any contrast’ is that any one who

said this would have *first* gone out of his way to find a form of words [‘but’] which introduced the implication, and *then* would have gone to some trouble to take it out again. Why didn’t he just leave it out?” (1961, pp. 137-138).

¹⁷After all, the semantics-pragmatics distinction is a tool of theoreticians. Chomsky puts it nicely: “We may make an intuitive judgment that some linguistic expression is odd or deviant. But we cannot in general know, pretheoretically, whether this deviance is a matter of syntax, semantics, pragmatics, belief, memory limitations, style, etc., or even whether these are appropriate categories for the interpretation of the judgment in question. It is an obvious and uncontroversial fact that informant judgments and other data do not fall neatly into clear categories: syntactic, semantic, etc.” (quoted in Schaffer 2004, p. 146).

¹⁸Nor is it being denied that there are other things we communicate by knowledge utterances (“I know that *p*” can be used to mean ‘You don’t have to tell me’ (Grice 1989, p. 53), e.g.), and other ordinary uses of knows (such as those wherein ‘knows’ does not take a propositional complement), which I have not discussed. This omission is harmless enough for present purposes, as they are not the sorts of examples which the revisionist arguments discussed here seek to exploit.

¹⁹Likewise, Ayer argues that the subject’s having ‘the right to be sure’ (1956, p. 35) (cf. (15)) is among the truth conditions of ‘*S* knows that *p*’ on the basis of examples such as this: “How does [the] man who knows what the results of the lottery will be differ from one who only makes a series of lucky guesses?...The difference is that to *say* that he knows is to concede him the right to be sure, while to *say* that he is only guessing is to withhold it” (1956, p. 33; emphases added).

²⁰As Bach says (forthcoming, Section 1), depending on the context, “John has beautiful handwriting” could be used to implicate that he’s poorly suited for a certain job, but also such things as that “[he] is a mediocre student, that he would make a fine translator, that he understood something he heard, or that he had no excuse for the sloppy paper he wrote.”

²¹Here, I am drawing in part upon Bach’s (2005, Section V) suggested means of handling Cohen’s airport

case, discussed below. As Bach notes, you cannot coherently attribute knowledge that *p* to another if you yourself are unsure as to *p*. What's added here is the suggestion that challenging another's having a right to be sure is one way in which one's doubt as to *p* can be expressed, along with the idea that by denying or suggesting the negation of a presupposition of a given assertion, "*q*", one can imply things about *q* itself.

²²As we'll see below, Unger takes his heavy use of stress to be legitimate; indeed, he believes that proper philosophical methodology recommends its employment in arguments such as his own. As we'll also see, however, his rationale for this policy in fact plays into a further anti-revisionist response to his arguments.

²³Insofar as the odd-sounding utterances Lewis cites are, as he says, merely explicit statements of the fallibilist idea, the latter portion thereof ('*S* cannot eliminate a certain possibility in which not-*p*,' e.g.) expresses the idea that the proposition in question is not certain, in the sense indicated above – i.e., that the subject's evidence doesn't entail what is (allegedly) known (and thus, the negation of any contrary propositions). According to Jason Stanley (2005), this is not the best reading of the truth conditions of the relevant sentences and, specifically, of the possibility clauses they contain. Responding in part to a previous presentation of the suggestion that the 'concessive knowledge attributions' (CKAs) Lewis cites are merely pragmatically inconsistent (Rysiew 2001, pp. 492ff.), Stanley argues that on the correct account of those possibility clauses, while such CKAs are self-contradictory, Lewis misses his intended target – the relevant sentences don't capture the fallibilist idea after all. Here, I assume with Lewis that the sentences in question express the fallibilist idea, but argue that his argument fails nonetheless. For a defence of this assumption against Stanley's arguments, and for further elaboration of the ideas of this paragraph, see Dougherty and Rysiew (forthcoming).

²⁴Compare Austin (1946, p. 98): "'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong,' just as you are prohibited from saying 'I promise I will, but I may fail'."

²⁵Sosa's proposal is related by Stewart Cohen (1999, p. 60).

²⁶Note once again the apparent contrast with familiar conversational implicatures, such as the aforementioned letter-writer case, wherein what's implicated does depend very much upon the specific context (see above, including n. 20).

²⁷Because, as both Bach (2005) and Brown (2006) note, self-attributions of knowledge may not be the best cases on which to focus (see too DeRose 2005, pp. 181-183), I consider here Cohen's third-personal airport example. (In my 2001 paper I focussed on DeRose's (1995) Bank Case. In part because of this, though its essential elements are retained here, the discussion of that paper was too geared towards the first-person case - - as John MacFarlane, e.g., has noted (2005, Section 3.1.1(b)).)

²⁸Especially in light of what's just been said about Unger's use of the device, one might worry about the stress on 'know' in Cohen's example. Not all such contextualist examples, however, involve the use of stress -- DeRose's (1995) Bank Case, e.g., does not (though it's tempting to suppose that it is slightly less natural as a result). But even absent such stress, there's reason to think that the truth-conditional content of what is said may not be to the fore in such cases (this Section); and suggestions to the contrary give rise (next Section) to a different sort of worry about the contextualist's argument.

²⁹Whether the latter result obtains depends upon the details of the example. DeRose insists that it is only in the right kind of case (or paired cases) that the contextualist really will claim that the relevant sentences are not conflicting -- i.e., the parties to the respective utterances must reasonably differ in their epistemic standards (most plausibly because of differences in their practical concerns), and there must be no "arm-twisting, arguing or convincing" between those parties, but only each of the two speakers' 'informing' their respective interlocutors, and meeting with no resistance to the claim, that the subject does/doesn't "know" (2005, pp. 172-181; forthcoming, Section 1). Cohen's airport case seems to satisfy such requirements.

³⁰DeRose for example allows, as a general point, that speakers can easily confuse 'warranted assertability' for truth (1999, p. 196) -- that there is "a general tendency to confuse what a speaker merely represents as being the case in making an assertion with what the speaker actually asserts" (2002, p. 186).

³¹See Rysiew (2001, 2005) and Brown (2006) for further elaboration and defence of the general strategy

employed in this paragraph for handling the sort of examples which contextualists cite. Rysiew (2005) and Brown (2006) include responses to DeRose's (1999, 2002) arguments for pessimism about the prospects of such an account in general; Rysiew (2005) includes a response to DeRose's (2002) criticism of the specific such account offered in Rysiew (2001).

Echoing one of DeRose's objections to the general strategy (see esp. DeRose 2002, pp. 192-193; 1999, p. 199), MacFarlane says: "although worries about misleading implicatures may be good reasons to refrain from asserting something, they aren't good reasons to assert its *negation*" (2005, Section 3.1.1(b)). The central idea here may be correct in cases where speakers are not committing any kind of error -- something that MacFarlane assumes for purposes of evaluating pragmatic theories. But, of course, that our pretheoretic intuitions as to what we're 'saying' are generally insensitive to the semantic/pragmatic distinction is essential to the view being presented here. And, to the extent that speakers *are* reading merely pragmatically-generated information onto the relevant sentences, it's not that they will utter a sentence that they know to be false (as, suppose, Mary and John's knowledge denial is) only so as to communicate some further, true information (e.g., that Smith doesn't meet certain unusually high epistemic standards); whether they see themselves as 'saying' something true/false will depend on the perceived truth value of the information that is merely pragmatically generated. (See Rysiew 2001, Sections 3-4, 10; 2005, p. 62; cf. Bach 1994, 2002; Brown 2006, Section 3.) Relatedly, to the extent that it is some merely pragmatically-generated information that a speaker intends to communicate, as opposed to (vs. in addition to) that expressed by the sentence itself, the case shouldn't be assimilated to the category of implicature (see Section 4, above). As MacFarlane notes (*ibid.*, Section 3.2), anyone who posits speaker error to explain away some linguistic uses must be on guard against eroding whatever evidential support other uses are supposed to provide for his/her own preferred semantic theory. But there's no reason to think that the process of distinguishing uses which are semantically significant from those which are not must be *ad hoc*. The present paper as a whole can be seen as one attempt to give principled reasons for adhering to the ho-hum view and the uses which support it while explaining away certain allegedly pro-revisionist linguistic data.

³²See Conee (2005, pp. 54-55) on the “delicate balance” contextualists must strike on this score.

³³Here, it’s worth noting that while Weinberg *et al.* (2001) have argued that epistemological intuitions are more variable and plastic than is commonly supposed, one such intuition that was found to be robust and univocal is that lucky guesses don’t count as knowledge. (*Ibid.*; and S. Nichols, personal communication.)

³⁴Perhaps, for reasons related to the principle of charity, one should generally attribute only beliefs which one regards as true, and perhaps justified as well. But it’s not clear how such a default norm would help the imagined theorist explain why, in third-person cases, we think that satisfaction of the ho-hum truth and/or justification conditions is required for knowing. For not only do we have a strong tendency to withdraw attributions of knowledge when it turns out that the subject’s belief is false or merely accidentally true, in those same cases we *don’t* withdraw attributions of belief, but in fact sometimes rely upon such attributions to do explanatory work, even citing their very falsity or unjustifiedness.

³⁵Very early versions of this paper were presented at Simon Fraser University and The University of Victoria. My thanks to audience members at those talks, as well as to Peter Graham, Melinda Hogan, Cindy Holder, Dom Lopes, and especially to Kent Bach, Mohan Matthen, and two anonymous referees for this Journal. Their thoughtful critical feedback at various stages helped to make the paper much better than it would otherwise have been.

REFERENCES

- Alston, William P. (2000) *Illocutionary Acts and Sentence Meaning* (Ithaca and London: Cornell University Press).
- Austin, J. L. (1946) “Other Minds,” *Proceedings of the Aristotelian Society*, Supplementary Volume XX. Reprinted in Austin *Philosophical Papers*, 3rd edition (New York: Oxford University Press, 1979): 76-116.
- Ayer, A. J. (1956) *The Problem of Knowledge* (Penguin Books).
- Bach, Kent (1994) “Conversational Implicature,” *Mind & Language* 9: 124-162.
- ___ (1999) “The Semantics-Pragmatics Distinction: What It Is and Why It Matters,” in Ken Turner, ed., *The*

Semantics-Pragmatics Interface from Different Points of View (Oxford: Elsevier): 65-84.

___ (2000) "Quantification, Qualification and Context: A Reply to Stanley and Szabó," *Mind & Language* 15.2-3: 262-283.

___ (2002) "Seemingly Semantic Intuitions," in J. Keim Campbell, M. O'Rourke, and D. Shier, eds., *Meaning and Truth* (New York: Seven Bridges Press): 21-33.

___ (2005) "The Emperor's New 'Knows,'" in Preyer and Peter, eds.: 51-89.

___ (forthcoming) "The Top 10 Misconceptions about Implicature," in Betty J. Birner and Gregory Ward, eds., *Drawing the Boundaries of Meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn* (John Benjamins): <http://userwww.sfsu.edu/~kbach/TopTen.pdf> (posting of Oct. 2006).

___ and Mike Harnish (1979) *Linguistic Communication and Speech Acts* (Cambridge, MA: The MIT Press).

Bonjour, Lawrence (2002) *Epistemology* (Lanham, Maryland: Rowman & Littlefield).

Brandom, Robert (1994) *Making it Explicit* (Cambridge: Harvard University Press).

___ (1995) "Knowledge and the Social Articulation of the Space of Reasons," *Philosophy and Phenomenological Research* 55: 895-908. Reprinted in Sosa and Kim, eds., *Epistemology: An Anthology* (Malden, MA: Blackwell, 2000): 424-432.

Brown, Jessica (2006) "Contextualism and Warranted Assertibility Manoeuvres," *Philosophical Studies* 130: 407-435.

Carston, Robyn (1988) "Implicature, Explicature, and Truth-theoretic Semantics," in *Mental Representations: The Interface Between Language and Reality*, edited by R. M. Kempson (Cambridge: Cambridge University Press): 155-181.

Cohen, Stewart (1999) "Contextualism, Skepticism, and The Structure of Reasons," *Philosophical Perspectives* 13: *Epistemology*, edited by James E. Tomberlin (Atascadero: Ridgeview): 57-89.

Conee, Earl (2005) "Contextualism Contested," in M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology* (Blackwell): 47-56.

Davis, Steve, ed. (1991) *Pragmatics: A Reader* (Oxford University Press).

DeRose, Keith (1992) "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52.4: 913-929.

___ (1995) "Solving the Sceptical Problem," *The Philosophical Review* 104.1: 1-52.

___ (1999) "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology*, Greco and Sosa (eds.) (Cambridge, MA: Blackwell): 187-205.

___ (2002) "Assertion, Knowledge and Context," *The Philosophical Review* 111.2: 167-203.

___ (2005) "The Ordinary Language Basis for Contextualism and the New Invariantism," *The Philosophical Quarterly* 55.219: 172-198.

___ (forthcoming) "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments Against Contextualism," *Philosophy and Phenomenological Research*:

<http://pantheon.yale.edu/%7Ekd47/Bamboozled.pdf> (posting of Sept. 2006).

Dougherty, Trent, and Patrick Rysiew (forthcoming) "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions," *Philosophy and Phenomenological Research*.

Dretske, Fred (2000) "Entitlement: Epistemic Rights Without Epistemic Duties," *Philosophy and Phenomenological Research* LX, 3: 591-606.

Feldman, Richard (2003) *Epistemology* (Upper Saddle River, NJ: Prentice-Hall).

Grice, H. P. (1961) "The Causal Theory of Perception," *The Aristotelian Society: Proceedings, Supplementary Volume* 35: 121-152.

___ (1989) *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).

Harnish, R. (1976) "Logical Form and Implicature," *An Integrated Theory of Linguistic Abilities*, (eds.) T.

Bever, J. Katz, T. Langendoen (Crowell). Reprinted in Davis, ed.

Hawthorne, John (2004) *Knowledge and Lotteries* (Oxford: Clarendon Press).

Kornblith, Hilary (2002) *Knowledge and Its Place in Nature* (New York: Oxford University Press).

Lewis, David (1996) "Elusive Knowledge," *Australasian Journal of Philosophy* 74.4: 549-567.

-
- Ludlow, Peter (2005) "Contextualism and the New Linguistic Turn in Epistemology," in Preyer and Peter, eds.: 11-50.
- MacFarlane, John (2005) "The Assessment Sensitivity of Knowledge Attributions," in Tamar Szabó Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology 1* (Oxford: Oxford University Press): 197-233.
- Malcolm, Norman (1986) *Wittgenstein: Nothing is Hidden* (Cambridge, MA: Blackwell).
- Moore, G. E. (1993) *Commonplace Book 1919-1953* (Bristol: Thoemmes Press).
- Neale, Stephen (1990) *Descriptions* (Cambridge, MA: The MIT Press).
- Plantinga, Alvin (1993) *Warrant and Proper Function* (New York: Oxford University Press).
- Plato (1981) *Five Dialogues*, translated by G. M. A. Grube (Indianapolis: Hackett).
- Preyer, Gerhard, and Georg Peter, eds. (2005) *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Clarendon Press).
- Quine, W.V., and J.S. Ullian (1978) *The Web of Belief*, 2nd edition (New York: Random House).
- Radford, Colin (1966) "Knowledge – By Examples," *Analysis* 27: 1-11. Reprinted in *Knowing: Essays in the Analysis of Knowledge*, 2nd edition, edited by Michael D. Roth and Leon Gallis (Lanham, MD: University Press of America, 1984): 171-185.
- Recanati, Francois (1986) "On Defining Communicative Intentions," *Mind & Language* 1.3: 213-42.
- ___ (1989) "The Pragmatics of What is Said," *Mind & Language* 4: 296-329.
- ___ (1993) *Direct Reference: From Language to Thought* (Cambridge, MA: Blackwell).
- Rysiew, Patrick (2001) "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35.4: 477-514
- ___ (2005) "Contesting Contextualism," *Grazer Philosophische Studien* 69: 51-70
- Sainsbury, Mark (1984) "Saying and Conveying," *Linguistics and Philosophy* 7.4: 415-432.
- Schaffer, Jonathan (2004) "Skepticism, Contextualism, and Discrimination," *Philosophy and Phenomenological Research* 69: 138-155.
- Searle, John R. (1969) *Speech Acts* (Cambridge, MA: Cambridge University Press).
- Sosa, Ernest (1991) *Knowledge in Perspective* (Cambridge, MA: Cambridge University Press).

Sperber, Dan, and Dierdre Wilson (1986) *Relevance: From Communication to Cognition* (Cambridge, MA: Blackwell).

Stalnaker, Robert C. (1974) "Pragmatic Presuppositions," in *Semantics and Philosophy*, edited by M.K. Munitz and P.K. Unger (New York University Press). Reprinted in Davis (ed.): 471-481.

Stanley, Jason (2005) "Fallibilism and Concessive Knowledge Attributions," *Analysis* 65.2: 126-31.

Unger, Peter (1975) *Ignorance* (Oxford: Clarendon Press).

___ (1984) *Philosophical Relativity* (University of Minnesota Press).

___ (1986) "The Cone Model of Knowledge," *Philosophical Topics* 14.1: 125-178.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich (2001) 'Normativity and Epistemic Intuitions,' *Philosophical Topics* 29.1-2: 429-460.

Weiner, Matt (2005) "Must We Know What We Say?," *The Philosophical Review* 114.2: 227-251.

___ (2006) "Are All Conversational Implicatures Cancellable?," *Analysis* 66.2: 127-130.

Williams, Michael (2001) *Problems of Knowledge* (Oxford & NY: Oxford).

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