

‘Knowledge’ and Pragmatics

PATRICK RYSIEW

1. Introduction

Linguistic phenomena have long served as an important source of data for epistemology. Considerations of ‘what we would say’, our ordinary knowledge-attributing habits and inclinations, etc., have helped to shape and constrain theories of knowledge, justification, and so on. But such considerations need to be handled with care, since our linguistic behavior and intuitions are shaped by a variety of factors. One of these is semantics, which has to do with linguistic items and the information they encode. Just as important, however, is pragmatics, which concerns the information arising from the tokening of such items – i.e., from *utterances*, or people’s *saying* things -- and the means by which such information is generated and recovered.¹

It’s often assumed that pragmatics has work to do only in special cases – for example, in understanding implicatures. Take Grice’s well-known letter of recommendation example (1989: 33). Asked to write a testimonial for a student applying for a philosophy job, S writes (only)

- (1) Mr. X’s command of English is excellent, and his attendance at tutorials has been regular.

Here, S damns with faint praise, getting it across that his opinion of Mr. X as a philosopher is rather low. Similarly, in Grice’s garage example (1989: 32), A is standing by an obviously immobilized car. B approaches, and the following exchange occurs:

- A: I’m out of petrol.
(2) B: There’s a garage around the corner.

Here, Grice says, B implicates that the garage is, or at least may be, open.

The need for pragmatics is hardly restricted to such examples, however -- and not just because pragmatic processes may also be essential to reference-fixing and resolving ambiguities. In standard examples of implicature, the speaker means what he says, but says it in order to communicate something further. Another sort of case that has gotten a lot of recent attention is a speaker’s meaning and communicating something *other than* what s/he says. Such common departures from literal speech don’t involve using words figuratively or ironically. Rather, they involve speakers’ meaning and communicating something *more or less closely related to* what they say -- something that “resembles” it in some contextually salient way (Sperber & Wilson 1986b). (So the relevant form of non-literality is ‘sentence nonliterality’, as Bach (2001) calls it – none of the elements is being used figuratively, but the speaker doesn’t mean quite what s/he says.) To take a well-worn example, an utterance of

¹ Theorists disagree as to how best to formulate the semantics-pragmatics distinction. The above characterization follows Bach (1999), who also argues against other formulations.

(3) ‘I’ve had [/haven’t had] breakfast’,

will typically communicate that the speaker has [/hasn’t] had breakfast *that day*. And an utterance of

(4) ‘You won’t die’,

said by a mother to her injured child, is liable to communicate that the child won’t die *from the cut* – hence, that the injury isn’t serious --, not that the child is immortal (Bach 1994).

As such easily multiplied examples show, linguistic content often underdetermines communicated content. Though united in wishing to counter “the compulsion to treat all pragmatically derived meaning as implicature” (Carston 1988: 176), theorists differ as to the proper taxonomic handling of cases like those involving (3) and (4). (Sperber & Wilson (1986a) and Carston (1988) call them ‘explicatures’; Recanati (1993), ‘strengthenings’ of sentence meanings; and Bach (1994), ‘implicatures’.) They also sometimes disagree about where specific examples fit within a given taxonomy,² how often sentences fail to express complete propositions, and whether ‘what is said’ is restricted to what’s encoded in the words uttered or includes the qualifying material likely communicated by utterances of (2) and (3), for example.³

Such disagreements shouldn’t obscure certain important points of agreement, however. First, all of the theorists just mentioned agree that what’s communicated by a given utterance often departs from the information encoded in the uttered words. Second, they agree that the information communicated by an utterance can be more or less closely related to what’s literally expressed. Thus, the content of sentence (1) differs markedly from what it’s used to implicate, whereas (3) and (4) and what they’re used to communicate are closely related.

What about (2)? Insofar as B communicates *that A can get service*, by the preceding reckoning, that would be counted as an implicature, since that information is hardly just an elaboration of the content of the uttered sentence. In this respect, the example resembles that involving (1). However, insofar as B’s utterance of (2) communicates *that there’s a garage around the corner that is (/might be) open* – a conceptual strengthening of the uttered sentence’s content -- it resembles the examples involving (3) and (4). Of course, there’s no tension in seeing the example as involving both of these facts: as implicatures illustrate, one can perform one speech act and thereby (indirectly) perform another.

² For example, Bach (2006: 28-29) suggests that most ‘scalar implicatures’ are really implicatures; and Bach (1994: Section 5) and Bach & Harnish (1979: Ch. 4) argue that metaphor and irony are best thought of as non-literal but direct speech acts, rather than implicatures.

³ Here, I use ‘what is said’ (as well as ‘linguistic content’ and ‘the content of the sentence’) to refer to the information encoded in expression-types, *modulo* disambiguation and reference-fixing.

In any case, and to mark a final point of convergence among the relevant theorists, it's generally agreed that, in spite of the noted differences among examples such as those above, fundamentally the same type of pragmatic processes are involved in all. Indeed, both 'neo-Griceans' such as Bach and 'post-Gricean' relevance theorists such as Sperber & Wilson and Robyn Carston hold that the same very general principles are in play in understanding *literal* utterances. For, as Bach puts it, that the speaker means what she says is itself something that needs to be inferred, and inferred on the same grounds as her meaning something additional to or other than what she says (Bach 2006: 24-25; cf. Sperber & Wilson 1987: 708).

What grounds are those? According to one well-established view, what enables speakers to communicate what they intend, even when this departs from or goes beyond what they say, is the fact that our conversational exchanges are governed by Grice's (1989) Co-operative Principle (CP).⁴ Or better, they are governed by the mutual presumption that others conform to CP (Bach & Harnish 1979: 62-65; Bach 2006: 24).

In discussing CP, Grice suggested various sub-maxims, under the headings of 'Quantity', 'Quality', 'Relation' and 'Manner'. Since it figures centrally in the discussion below, the role of relevance ('Relation') should be briefly illustrated. Thus, it's because B, in uttering (2), is presumed to be intending something that's relevant that his saying what he does communicates that the nearby garage is (/might be) open. Similarly, in the examples involving (3) and (4), it's not the obvious falsity of what's said -- e.g., "You're not going to die" -- that triggers the inference to the information the speaker intends, any more than the obvious truth of the negated claim ("You're going to die") would be what the speaker relied upon in communicating that the injury was extremely serious. In either case, rather, it is a lack of relevant specificity in what is said that invites the appropriate inference (Bach 2001: 255-256).

A further crucial point illustrated by these examples is that conformity to CP doesn't require that the sentences uttered be maximally relevantly informative.⁵ Indeed, as Saul puts it, "a key way of generating implicatures, according to Grice, is to violate maxims at the level of what is said" (2002: 364). Thus, for example, in their respective contexts, sentence (1) is under-informative, and (2) is not sufficiently relevant. In such cases, it's because the uttered sentences are not themselves maximally relevantly informative that speakers are able to communicate things over and above, or other than, what they say: they violate CP at the level of what is said but -- and thereby -- conform at level of what's communicated (see too Neale 1990: 106, n. 19; Grice 1989: 32-35; van der Henst *et al.* 2002: 458).

One prominent alternative to the Gricean framework is relevance theory. According to Sperber & Wilson's *principle of relevance*, "every act of ostensive communication communicates the presumption of its own optimal relevance" (1986a: 158). Optimal relevance, in turn, is understood to be a matter (roughly) of minimizing processing costs

⁴ "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" (1989: 26).

⁵ The latter is Harnish's (1976) suggested compression of CP.

while maximizing contextually relevant informational payoff (Sperber & Wilson 1986a). There are important questions about the extent to which Gricean and relevance-theoretic approaches are compatible (e.g., Saul 2002, Bach 2010). In what follows, however, we'll tend to employ the Gricean framework, though the key points could likely be rendered just as well in relevance-theoretic terms.⁶

2. Pragmatics and contextualism

It's uncontroversial among epistemologists that some data call for a pragmatic handling. For example, consider apparently non-factive uses of 'know(s)', as when it is said that ancient thinkers "knew" that the earth was flat. The majority view is that such utterances should not be taken literally – more specifically, that they involve 'protagonist projection' (Holton 1997), reflecting the point of view of the subject(s). Another example is Moore-paradoxical assertions – i.e., utterances of the form "*p*, but I don't believe that *p*" (Moore 1942: 543). These are odd-sounding, but not for semantic reasons. Neither must they be seen as involving implicature or implicature (/explicature), either of which would require that the speaker *intend* to communicate (with the utterance of the first half of the sentence, "*p*") that they believe that *p*. Instead, the oddity can be explained by the familiar fact that belief is a condition on felicitous assertion (Grice 1989: 42; Bach 2006: 26) and thus that, whether or not one so intends, in asserting one represents oneself as believing (Rysiew 2007). Because that's so, and because pragmatically generated information can influence our intuitive response to an utterance, Moore-style assertions can persist in causing discomfort even in those who are aware that they involve no semantic inconsistency.

More controversial is whether pragmatics can help⁷ explain the data that figure centrally in arguments for epistemic contextualism (EC), such as our intuitive responses to DeRose's Bank (1992) and Cohen's Airport (1999) Cases. In such examples, a scenario ('Low') in which a knowledge attribution seems correct is paired with one ('High') in which denying knowledge of the same subject seems correct, even though none of what would normally be regarded as epistemologically relevant factors have changed. The only clear difference between the cases is an increase in the practical importance of the subject's getting it right. EC is said to provide a natural resolution of this apparent tension: the content of sentences of the form '*S* does/doesn't know that *p*' vary, depending on the context, where 'context' is take to refer to features of the knowledge attributor(s)' psychology and/or conversational-practical situation. Acontextually, a knowledge sentence expresses (something like),

- (5) *S* has a true belief and is in a strong_C epistemic position with respect to *p*,

where strength of epistemic position is a function of the standards operative in the context, *C*; or, more simply,

⁶ For a particular application of relevance theory to the case of EC, see Jary & Stainton, ch. 37.

⁷ There is no need to claim exclusivity here. Pragmatic explanations can complement psychological ones, for example. After all, pragmatics is already as much about psychology – the psychology of utterance-production and -interpretation – as anything.

- (6) *S* has a true belief that belief satisfies the contextually operative epistemic standards.⁸

With this view in place, the apparent tension between the relevant (Low) attribution and (High) denial dissolves. For what the relevant sentences actually express are,

(Low) *S* has a true belief and is in a *moderately* strong epistemic position with respect to *p*, and

(High) It's not the case that (*S* has a true belief and is in a *very* strong epistemic position with respect to *p*).

And these, of course, are perfectly compatible.

While EC has the result that each of the paired claims is true, a pragmatic response proposes that one member of the pair is false, even if the relevant assertion is pragmatically appropriate, and even if it seems true. (Following DeRose 1999, it has become common to refer to such views as involving a 'warranted assertability manoeuvre', or 'WAM'.) Some such accounts are sceptic-friendly, holding that the attribution of knowledge in Low is false, even if appropriately made (e.g., Unger 1975, Blauw 2003, Schaffer 2004). Other, moderate invariantist views (e.g., Rysiew 2001, 2005, 2007, Black 2005, Brown 2006, Hazlett 2007, and Pritchard 2010) explain, in pragmatic terms, why the knowledge-denial in High is apt, and even seems true, in spite of its literal falsity. There is disagreement among these theorists as to the proper form such an explanation should take. Below, we'll consider an account that appeals to Grice's maxim of Relation ('Be relevant!').

DeRose himself is sceptical about the prospects for a successful WAM here. In plausible WAMs,⁹ DeRose says, we have some genuinely conflicting intuitions, such that we know in advance that something's going to have to be explained away. Further, a good WAM explains away one of the conflicting intuitions by appealing to the fact that the utterance of some truth generates a false implicature. (And it does so by appeal to some very general conversational principles.) But in the Bank and Airport Cases, these 'good-making' features are absent: First, our intra-contextual intuitions are in harmony (2009: 116). Second, and relatedly, the invariantist will have to explain an apparent truth in terms of an utterance of some *falsehood* generating a *true* implicature. And this, DeRose thinks, is implausible: "For, except where we engage in special practices of misdirection, like irony or hyperbole, don't we want to avoid falsehood both in what we implicate and (especially!) in what we actually say?" (*ibid.*: 114).

⁸ "When the content of knowledge-attributing sentences varies from context to context, what is varying are the epistemic standards, or how strong an epistemic position the subject must be in to count as knowing; the content of a given use of 'S knows that p' is that S (has a true belief that p and) meets the epistemic standards relevant to the context of utterance" (DeRose 2009: 34).

⁹ DeRose (2009: Chapter 3) illustrates with the example of possibility statements, executing a 'WAM' on behalf of his preferred account of the semantics thereof.

As Brown (2006: 411-419) -- and, following her, Rysiew (2005: 59-62) -- have argued, however, DeRose's scepticism about the prospects for a pragmatic handling of the relevant data is unwarranted. First, if the examples in question did not give rise to any apparent inconsistency -- when we compare our intuitions across the Bank and Airport cases, for example -- there would be no reason to consider them interesting, much less to think that they might call for a contextualist handling. EC might imply that such an inconsistency is merely apparent, but that there is an inconsistency is an apparent truth all the same.

As to DeRose's second point, it may be true -- special cases of indirection aside -- that we want to avoid falsehood in what we say. But that's only insofar as we think that what we say is what we mean. However, as explained above, it's an established view that in many cases not involving any kind of irony, hyperbole, or figurative speech, the content of the uttered sentence fails to correspond to what the speaker intends to communicate. On that view, irony, figurative speech, etc., are limiting cases on a continuum of less than fully literal speech, rather than the best or only examples thereof. Further, as Brown says, "there is a wide range of non-figurative uses of language where one standard approach involves supposing that a literally false statement seems correct because it pragmatically conveys a truth" (2006: 415).¹⁰ Indeed, that a speaker might make a literally false statement to communicate some truth is a corollary of the fact that violations of CP at the level of what is said are allowable when they enable conformity at the level of what's communicated. And that, in such a case, the statement itself can seem true is a consequence of the fact -- which DeRose's own favored WAM exploits -- 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" (2009: 106; cf. 83).

With the previous points in place, the issue is simply whether there exists a plausible pragmatic explanation of the relevant data in terms of some general, and generally-accepted, conversational principles. According to Rysiew and Brown, Gricean considerations operating through the maxim of Relation provide such an account. By way of background to her own presentation of such a view, Brown (2006: 425) illustrates how considerations of relevance function in Grice's garage example, introduced above. We'll do that here as well, in part because it will give us occasion to address a worry about the example that also arises, in general form, in cases of central interest to EC.

So, again: given A's evident need for gasoline, B's utterance of

(2) There's a garage around the corner,

implicates that the garage is or might be open; for otherwise the utterance would violate Relation. (Again, we may think this is an implicature; but either way, what's communicated is something that closely resembles what's said; and either way, that it's

¹⁰ The experiments of Gibbs & Moise (1997), Nicolle & Clark (1999), and Novacek (2001), e.g., strongly suggest that intuitive assessments of 'what is said' are influenced by pragmatic factors. Indeed, Nicolle & Clark report that, in some cases, when asked to select the paraphrase that best reflects what a speaker has said subjects choose clear cases of implicatures.

communicated is explained by the same general principles.) By the same token, if B thought that the nearby garage wasn't (at least possibly) open, his utterance would be misleading. "Rather," Brown continues, "in such a scenario, and assuming that she believes that there is no other nearby open garage, it would be correct for her to reply 'No, there's no garage nearby'. While this utterance is literally false, it pragmatically conveys the true claim that there is no open garage nearby" (*ibid.*).

According to Blome-Tillmann, however, Brown's handling of the case fails. As Blome-Tillmann sees it, Brown is here assuming that, at least in the case of implicatures driven by Relation, the following principle holds:

Converse Implicatures (CI):

If an utterance of a sentence S conversationally implicates p in C, then an utterance of not-S conversationally implicates not-p in C. (Blome-Tillmann 2013: 4306)

However, CI is false of many implicatures, including many Relation implicatures.¹¹ Further, while agreeing that "B's utterance [here] conveys that there is no open garage nearby," Blome-Tillmann denies that this is an implicature. Rather, B's utterance communicates the latter because the proposition that there's no garage nearby *entails* that there's no open garage nearby (*ibid.*: 4307, n. 32). He continues:

The fact that the phenomenon at issue is not a conversational implicature can be demonstrated further by the cancellability test: 'There's no garage nearby; but there's an open garage nearby' is contradictory in all contexts, because it semantically expresses a contradiction. (*Ibid.*)

In response, note first that Brown's handling of the example is hardly idiosyncratic: Garcia-Carpintero (2001: 113), e.g., cites the case as illustrating the inheritance of implicated contents in more complex constructions. Second, contrary to what Blome-Tillmann implies, one can implicate something that's entailed by what one says (Davis 2014: Section 4; Bach 2006: 24). And of course, where what's implicated is entailed by what one says, cancellability will not be an appropriate test. Third (and again contrary to Blome-Tillmann), that the proposition that there's no garage nearby entails that there's no open garage nearby isn't sufficient for the latter's being communicated: not everything entailed by what one says is communicated; for something to be communicated, it must be plausibly regarded as intended (meant) by the speaker.¹² So we'd still need an account

¹¹ Blome-Tillman cites (*ibid.*: 4307, n. 32), as a Relation-based counterexample to CI, Grice's letter-writer case – i.e., the example involving (1) above. 'Praising with faint damns' is a real phenomenon. But set that aside – the discussion here doesn't require that CI be true, even that it be true of all Relation implicatures. All that's required is (a) that 'pragmatic inheritance' does occur, and (b) that it is plausibly thought to occur in the present case (as well as in the cases of concern to EC). As to (a), implicatures [/explicatures] and definite descriptions provide examples (Brown 2006), as do (according to some) belief reports (Berg 2012 and others). As to (b), see the discussion below.

¹² And intended to be communicated in part by means of the recognition of that intention. That communication involves such reflexive intentions is a key insight of Gricean theory.

of why this particular entailment *is* communicated; and again, Relation provides the answer.¹³

Finally, consider the following natural elaboration of the case:

A: I'm out of petrol.

B: I'm sorry, there's no station nearby.

C (onlooker, to B): Isn't there a station just two blocks over?

B: Yes, but it's been closed for months.

Here, the two things B says can't both be true. Nonetheless, with his second utterance B is most naturally read, not as contradicting himself, but as making it clear just what he meant all along. For the same reason, if, after B has said, "There's no station nearby", A discovers that there is a nearby station but that it's shut down, there's no reason to think she would feel as though she'd been misled by B or that B had 'gotten it wrong'.

In sum, Blome-Tillmann has given us no reason to question Brown's handling of the garage example: her treatment of it faithfully illustrates the operation of considerations of relevance in everyday linguistic communication.

Let's proceed, then, to a specific proposed pragmatic handling of the data thought to lend support to EC. Briefly sketched, the account runs as follows: According to a widely-held view, knowing requires (unGettiered) true belief plus the subject's being in a good epistemic position with respect to the proposition in question. What makes such a view invariantist is the idea that there is a fixed standard for the level of epistemic goodness knowledge requires. Even so, because a speaker is presumed to be conforming to Relation, in uttering a sentence of the form, '*S* knows that *p*', *s/he* is naturally taken to intend/mean that *S*'s epistemic position with respect to *p* is 'good enough' *given the epistemic standards that are operative in the context in question*. For, in naturally occurring situations,¹⁴ it is only if speakers are so regarded that the presumption that their conversational contributions are maximally relevantly informative, and thus that they conform to CP, is preserved.

In the Airport Case, for example, it's mutually obvious that Mary and John want to ensure that their epistemic position with respect to the flight plan is *very* strong -- strong enough to rule out the possibility of a misprint in their itinerary. Being in an epistemic position of such strength may or may not be required for knowing. Either way, due to considerations of relevance, it would be misleading of Mary and John to attribute knowledge to Smith, and so represent him as being in a good epistemic position, if they thought that his epistemic position wasn't so good as to put their shared and pressing

¹³ Recall the discussion of utterances of (3) and (4) and their negations: the inference to what the speaker means is driven by Relation, not Quality ('Do not say what you believe to be false'); the presumed or likely truth-values of the sentences uttered – hence, what they do/don't entail – play no direct role.

¹⁴ As opposed to, say, in an epistemology class. There, if conditions are right, what attributions or denials of knowledge communicate might be just what the relevant sentences express.

concerns to rest. -- Just as, in the garage example, it would be misleading to utter (2), ‘There’s a garage nearby’, unless one thought it was (possibly) *open*. Parallel considerations apply to the knowledge denial: as *per* Brown above, the speaker who utters the negation of (2) is most naturally taken to mean that there is no *open* garage nearby; similarly, given that what’s relevant to Mary and John is whether Smith is in a *very* strong epistemic position, their utterance of “Smith doesn’t know...” is most naturally taken to mean that it’s not the case that Smith (has a true belief and) is in a *very* strong epistemic position – that his epistemic position isn’t so good as to put their shared and pressing concerns to rest, that he can’t rule out the possibility of a misprint, and so on.¹⁵ Given what we’re told, the latter thoughts seem not just highly relevant but true. And if they (we) read what’s conveyed by the relevant utterance onto the sentence uttered, the knowledge denial will strike them (us) as true as well.

3. Some objections considered

One objection to the account just sketched, in certain of its forms, is that it is *ad hoc*. In the language of Rysiew 2001, knowledge requires an ability to rule out the relevant not-*p* alternatives; but, thanks to considerations of relevance, an attribution of knowledge pragmatically conveys that *S* can rule out the *salient* alternatives – i.e., the not-*p* alternatives that the conversants have in mind.¹⁶ As DeRose at one time interpreted it (2002: 198, n. 17), Rysiew’s view is that “an assertion of ‘S knows that P’ carries two separate but related meanings” – a ‘semantic meaning’ and ‘a pragmatic meaning’; and the role of relevance is simply to get the hearer “to fasten on the second, pragmatic meaning” in cases where “the semantic content of ‘S knows/doesn’t know that P’ is clearly irrelevant to the purposes at hand” (*ibid.*; cf. 2009: 118ff.). As Rysiew (2005: 63-64) explains, however, on his view nothing ‘has two meanings’, and the role of relevance is (therefore) not simply to get the hearer to switch from one to the other. Rather, there is what the sentence means, and what the speaker means in uttering it; and Relation explains how an utterance of the former, given what it means, communicates the latter.

A second, more influential objection has it that the relevant pragmatic conveyances fail Grice’s cancellability test (Grice 1989: 44). On the proposed account, an utterance of ‘*S* knows that *p*’ might pragmatically convey that *S* is in a very strong epistemic position -- that *S* can rule out the not-*p* alternatives that are in play, that (consequently) there’s no need for further investigation, and so on. However, the objection runs, “*S* knows, but *S* is not in an especially strong epistemic position”, or “*S* knows, but *S* needs to check further,” sound odd (Cohen 1999: 60). In response, we should first note, again, that “[s]peech act content is, in the jargon, a massive interaction effect” (Borg 2012: 15) – i.e., that what’s communicated by a given utterance is the product of multiple factors (syntax, semantics, pragmatics, etc.). And, to the extent that our intuitions are attuned to the total message communicated, rather than just the semantic features of what is said, certain cancellations

¹⁵ There is no need to single out one such proposition as *the* thing the speakers mean. What’s communicated can be vague or indeterminate.

¹⁶ There are important issues here about the role that salience considerations play, how they interact with practical interests, and so on. For criticism of Rysiew (2001) on this, see Brown (2006); Rysiew (2007) is meant to incorporate the relevant points.

are apt to feel uncomfortable (Rysiew 2001: Section 7; Brown 2006: 428; cf. Chomsky 1977: 4).

Our general tendency to read merely pragmatically generated information onto what we say aside, however, is there any reason to think that we might make such an error in the case at hand? After all, to return to the letter-writer case discussed above, no one is going to confuse what's said with what's communicated *there*; so why think it happens here? Several things are worth noting. First, the pragmatic phenomena we're now considering are not 'particularized' (Grice 1989: 37): that the speaker takes *S* to satisfy the contextually operative standards, to be able rule out the salient not-*p* alternatives, and so on, is inferable simply from her asserting the relevant sentence and the presumption that she's conforming to CP; unlike the example involving (1), very little in the way of specific knowledge of the context in which the relevant sentences are uttered is required (see Rysiew 2007). Second, while in the case at hand what the speaker means isn't what s/he says, unlike familiar tropes, say, none of the elements of the knowledge sentence is being used non-literally; the example thus involves a kind of non-literality that's harder to spot (Bach 2001: 249-250).¹⁷ Third, and again unlike the letter-writer case, what's communicated in the present example closely resembles what's semantically expressed. Finally, while error about certain items in the lexicon might antecedently be rather unlikely, even those whose job it is to provide an account of the semantics of 'know(s)' disagree on the matter, sometimes significantly, and with some holding that the sort of thing said on the present account to be merely pragmatically conveyed is in fact required for knowing. For this reason too, it's not implausible that we might be prone to reading the latter type of information onto knowledge sentences.

Still, that we naturally make the mistake doesn't mean that we must do so. Thus, when the relevant information is spelled out explicitly, we can distinguish between (for example) a speaker's being in a strong, versus a *very* strong, epistemic position with regard to some *p* (the former being what knowledge requires, the latter being what an attribution of knowledge might convey). In fact, and returning to the issue of cancellability, that's essentially what's happening with certain instances of the allegedly problematic utterances, instances that *don't* sound odd at all. For instance, "*S* knows that *p*, but *S* can't rule out some *bizarre* not-*p* possibilities"; or, "*S*'s epistemic position with respect to *p* isn't so strong that her belief that *p* would match the facts in *very* distant not-*p* worlds; but still, she knows that *p*" (Rysiew 2001: 495; Brown 2006: 428).¹⁸ With these

¹⁷ Blome-Tillmann observes that even highly formulaic and idiomatic instances of hyperbole ('apologise a thousand times', 'cry a flood of tears') are easily detected as non-literal, as are metaphor and irony ('She's made of stone, this girl', A: 'John's an atheist' – B: 'Yes, and so is the pope') (2013: 4300-4302). Note that many such examples involve figurative speech, rather than the sort of non-literality being considered here. They also all typically exploit the obvious flouting of Quality (see n. 13) at the level of what is said -- in Blome-Tillmann's terms, such "utterances...trigger an implicature *because* their semantic content cannot be readily added to the common ground" (*ibid.*: 4310). This is not true of the case under consideration here.

¹⁸ Such statements are examples of *concessive knowledge attributions* (CKAs) (Rysiew 2001) – i.e., utterances of the form "*S* knows that *p*, but it's possible that *q*" (where *q* entails not-*p*). The proper handling of CKAs is another front on which contextualists and pragmatic theorists (among

more natural utterances, the speaker's presenting information in a certain order, using stress selectively, and so on, makes it easier to understand what he's trying to get across; and when that's so, the appearance of inconsistency recedes.

DeRose's most recent objection to the pragmatic account we're considering, however, takes a different form. Here, the worry is not that it is *ad hoc*, that it leads to violations of cancellability, or that it posits an implausible form of error. Rather, it's that considerations of relevance simply wouldn't operate in the way the account requires. Here is the objection, as applied to the High variant of the Bank Case:

“...how could relevance concerns drive the hearer from the semantic content posited by the theory to the theory's proposed pragmatic destination? The falsehood that, according to the account, the speaker has asserted would be very relevant to the concerns of the participants in this conversation. We will grant that the *issue* or the question of whether or not the subject knows_H^[19] is the most relevant issue to the purposes at hand (and is therefore more relevant than is the question of whether he knows_M)....But that the speaker doesn't even know_M that the bank is open on Saturdays – which according to Rysiew's account is what the speaker has asserted – would of course settle (in the negative) the salient question or issue of whether he knows_H. So, in saying that he doesn't 'know' that the bank is open on Saturdays, on Rysiew's account, the speaker asserts what would be an extremely relevant thought.” (DeRose 2009: 122)

Thus, DeRose concludes, “it's hard to see how relevance concerns would drive the hearer to suppose that anything other than what was said is what the speaker intends to communicate” (2009: 124).

This objection of course recalls Blome-Tillmann's criticism of Brown's handling of the garage example. The core idea of each is that all of the work is being done by the fact that the sentence uttered entails information that's most relevant to the subjects. In Blome-Tillmann's criticism, this was taken to show that what's entailed is communicated but not implicated; here, it's being used to argue that what does the entailing is most plausibly thought to be what the speaker means.

In the present case, one response is that it's questionable whether the relevant entailment

others) compete. Lewis (1996) argues that CKAs state the basic fallibilist idea but are self-contradictory, the best fix being to adopt a certain version of EC. Stanley (2005) denies that the relevant claims express fallibilist thoughts. Dougherty and Rysiew (2009) argue that they do, but that their (sometimes) infelicity is pragmatic.

¹⁹ 'Knows_M' and 'knows_H' are meant by DeRose to express *S*'s satisfying moderate/high standards for knowing, the idea being that, according to the moderate invariantist a knowledge attribution/denial that semantically encodes that *S* does/doesn't know_M that *p* might pragmatically communicate that *S* does/doesn't know_H that *p*. While the expressions are intended merely as shorthand, they are misleading, as the invariantist need hardly grant that there are degrees or types of knowledge.

holds.²⁰ In the language of Rysiew 2001, DeRose's objection assumes that the set of relevant alternatives is a subset of the salient alternatives. In more general terms, it assumes that strength of epistemic position is a straightforward matter of variation along some one (or several) neatly ordered dimensions (cf. Bach 2005: 57, n. 11; DeRose 2009: 7, n. 3). In either form, the assumption is dubious.

However, even granting the assumption – i.e., even granting that ‘*S* doesn't know_M that *p*' entails that ‘*S* doesn't know_H that *p*' (and setting aside concerns about the shorthand: see n. 19) -- the objection fails. For, first, both Grice and Sperber & Wilson make it clear that the fact that some information is somewhat relevant, even very relevant, won't suffice for its being communicated, and/or its being appropriately thought to be what the speaker intends.²¹ What must be presumed is that what the speaker intends, given what he's said, affords maximal relevant informational payoff for least interpretative effort. Nor, for the same reason, does it matter that because ‘*S* doesn't know_M that *p*' is (we are for the sake of argument supposing) logically stronger, it “give[s] more information” than ‘*S* doesn't know_H that *p*' (see DeRose 2009: 123-124); for that extra information might tax attentional and inferential resources without itself contributing to a resolution of the issue of greatest immediate interest.

This leads to the second point: the reason the speaker communicates that *S* doesn't “know_H that *p*” is the same reason that *B*'s utterance in the garage example communicates what *it* does – namely, because this further information speaks directly to “the most relevant issue to the purposes at hand” (*ibid.*: 122), and because it is only on the supposition that he intends to communicate that information, and expects the hearer to recognize that intention, that the presumption that he is conforming to CP is preserved. As Sperber & Wilson put it, “when the speaker could not have expected his utterance to be relevant to the hearer without intending him to derive some specific contextual implication from it, then...that implication is also an implicature” (1981: 284; cf. Bach & Harnish 1979: 169).

Finally, consider the following natural elaboration of the case:

- A: “Banks do change their hours. Do you know the bank will be open tomorrow [Saturday]?”
 B: “Well, no. I'd better go in and make sure.”
 C (onlooker, to B): “But wait. You were there just two weeks ago on a Saturday. You have no reason to think that its hours have changed. In fact, you have, overall, very good reason to believe that the bank is open tomorrow.”

²⁰ This is noted by Blome-Tillmann (2013: 4309-4310). Blome-Tillmann has his own concerns about the pragmatic account we're considering, only some of which are touched on here.

²¹ Sperber & Wilson (1986a) speak of *optimal* relevance. Bach (2010: 130; cf. Bach and Harnish 1979: 92) prefers *sufficient* relevance, which also permits the point just made. And Grice, in explicating Relation, writes: “I expect a [conversational] partner's contribution to be appropriate to the immediate needs at each stage of the transaction; if I am mixing ingredients for a cake, I do not expect to be handed a good book, or even an oven cloth (though this might be an appropriate contribution at a later stage)” (Grice 1989: 28).

B: “Yes. But it’s *very* important we get the check cashed before Monday. And I can’t be sure that it hasn’t changed its hours (etc.)”

Here, if a moderate invariantist view is correct, the two things B says can’t both be true (assuming his belief is true and unGettiered), since (given a true, unGettiered belief) being in an epistemic position as good as his apparently is suffices for knowing. Nonetheless, with his second utterance B is most naturally read, not as contradicting himself, but as making it clear just what he meant all along. For the same reason, if, after the fact, A and B learn that B did satisfy the (moderate) conditions that are, in fact, required for knowing, there’s no reason to think A would feel as though, in saying “No [I don’t know]”, B had misled her or had gotten it wrong.

In sum, then, even assuming that the relevant entailment holds, it is not “hard to see how relevance concerns would drive the hearer to suppose that anything other than what was said is what the speaker intends to communicate” (DeRose 2009: 124). Quite the contrary: the present reflections reinforce the idea that considerations of relevance favor the view that what the speaker intends to communicate is that his epistemic position isn’t so strong as to satisfy the elevated standards that are in play.

The latter, of course, is just what DeRose (and others contextualists) think the speaker here would communicate. However, contextualists themselves owe us an account of how that happens. As we saw at the outset, pragmatics are required even in cases where the speaker means just what he says. That what he says, according to EC, itself encodes the relevant context-relative information – viz., (High) It’s not the case that (*S* has a true belief and is in a *very* strong epistemic position with respect to *p*) -- does not explain how the hearer grasps that that’s what the speaker means (see Rysiew 2001: 506).²² And now notice that, if DeRose’s argument above is effective, it frustrates the contextualist at just this point. For if DeRose is correct, that *S* doesn’t ‘know_M’ that *p* suffices to address “the most relevant issue to the purposes at hand” (by settling it in the negative), and so there’s no reason for the hearer not to take *that* to be what the speaker means. The denial in High would then be correctly read as a simple negation of the Low attribution – which is precisely what contextualists want to deny. On the other hand, if there *is* a good account of how the hearer gets from an utterance of the relevant sentence to the speaker’s meaning that *S* “doesn’t know_H that *p*” (i.e., High), the pragmatic invariantist can avail herself of that as well. In short, both parties to the debate require a way of moving from the sentence uttered to what the speaker means thereby. Either there’s a plausible story to be told here, or there is not. If not, that’s equally problematic for both parties. But if so, it’s not clear what benefit accrues from taking the speaker’s meaning to be reflected in the content of the sentence itself (Rysiew 2001, Section 10).

4. Conclusion

The aim of this Chapter has been to provide an introduction to pragmatics and its relevance to debates surrounding EC. To give a good sense of how the relevant issues play out, we’ve focused on one particular pragmatic approach and some of the more prominent objections to it. The reader is reminded that there are other pragmatic accounts

²² Neale (2007: 81-82) makes a similar point.

(and, no doubt, other objections). Also, while we have focused here on the handling of Low-High paired cases, there are other fronts on which EC and pragmatic considerations compete. These include concessive knowledge attributions (see n. 18) and considerations of the social function of ‘know(s)’.²³ Like pragmatics itself, its bearing upon EC and the issues it brings to the fore is multi-faceted and under-appreciated.²⁴

References

- Bach, K. (1994) “Conversational Implicature”, *Mind & Language* 9: 124-162.
- ____ (1999) “The Semantics-Pragmatics Distinction: What it is and Why it Matters”, in K. Turner (ed.), *The Semantics-Pragmatics Interface from Different Points of View*. Oxford: Elsevier, pp. 65-84.
- ____ (2001) “Speaking Loosely: Sentence Nonliterality”, in P. French and H. Wettstein (eds.), *Midwest Studies in Philosophy* 25. Malden MA: Blackwell Publishers, pp. 249-263.
- ____ (2005) “The Emperor’s New ‘Knows’”, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Clarendon Press, pp. 51-89.
- ____ (2006) “The Top 10 Misconceptions about Implicature”, in B.J. Birner and G. Ward (eds.), *Drawing the Boundaries of Meaning*. Amsterdam and Philadelphia: John Benjamins, pp. 21-30.
- ____ (2010) “Implicature vs. Explicature: What’s the Difference?”, in B. Soria and E. Romero (eds.), *Explicit Communication*. New York: Palgrave MacMillan, pp. 126-137.
- ____, and M. Harnish (1979) *Linguistic Communication and Speech Acts*. Cambridge MA: The MIT Press.
- Berg, J. (2012) *Direct Belief*. Berlin: Walter de Gruyter.
- Blaauw, M. (2003) “WAMing Away at Contextualism”, *Nordic Journal of Philosophy* 4(1): 88–97.
- Black, T. (2005) “Classic Invariantism, Relevance, and Warranted Assertability Manoeuvres”, *The Philosophical Quarterly* 55(219): 328-336.
- Blome-Tillman, M. (2013) “Knowledge and Implicatures”, *Synthese* 190(18): 4293-4319.

²³ Henderson (2009) argues that such considerations provide a new rationale for EC. For a treatment of the topic in which pragmatics figure prominently, see Rysiew (2012).

²⁴ For valuable feedback, and for help at various stages with the ideas expressed herein, I’m indebted to Jonathan Adler, Kent Bach, Jessica Brown, Nathan Cockram, Jonathan Ichikawa, Mike Raven, and Rob Stainton.

- Borg, E. (2012) *Pursuing Meaning*. Oxford: Oxford University Press.
- Brown, J. (2006) "Contextualism and Warranted Assertibility Manoeuvres", *Philosophical Studies* 130: 407-435.
- Carston, R. (1988) "Implicature, Explicature, and Truth-theoretic Semantics," in R. M. Kempson (ed.), *Mental Representations*. Cambridge: Cambridge University Press, pp. 155-181.
- Chomsky, N. (1977) *Essays on Forms and Interpretation*. Amsterdam: North Holland.
- Cohen, S. (1999) "Contextualism, Skepticism, and The Structure of Reasons", in J.E. Tomberlin (ed.), *Philosophical Perspectives 13: Epistemology*. Atascadero CA: Ridgeview, pp. 57-89.
- Davis, W.A. (2014) "Implicature", *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2014/entries/implicature/>.
- DeRose, K. (1992) "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research* 52(4): 913-929.
- ___ (1999) "Contextualism: An Explanation and Defense", in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*. Cambridge, MA: Blackwell, pp. 185-203.
- ___ (2002) "Assertion, Knowledge and Context", *The Philosophical Review* 111(2): 167-203.
- ___ (2009) *The Case for Contextualism*. Oxford: Oxford University Press.
- Dougherty, T., and Rysiew, P. (2009) "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions", *Philosophy and Phenomenological Research* 78(1), 123-132.
- Garcia-Carpintero, M. (2001) "Gricean Rational Reconstructions and the Semantics/Pragmatics Distinction", *Synthese* 128: 93-131.
- Gibbs, R.W., and J. F. Moise (1997) "Pragmatics in Understanding What is Said", *Cognition* 62: 51-74.
- Grice, H. P. (1989) *Studies in the Way of Words*. Cambridge MA: Harvard University Press.
- Harnish, R. (1976) "Logical Form and Implicature", in T. Bever, J. Katz, and T. Langendoen (eds.), *An Integrated Theory of Linguistic Abilities*. New York: Thomas Y. Crowell, pp. 464-79.
- Hazlett, A. (2007) "Grice's Razor", *Metaphilosophy* 38(5): 669-690.

Henderson, D. (2009) "Motivated Contextualism," *Philosophical Studies* 142(1): 119-31.

Holton, R. (1997) "Some Telling Examples", *Journal of Pragmatics* 28: 625-628.

Lewis, D. (1996) "Elusive Knowledge", *Australasian Journal of Philosophy* 74(4): 549-567.

Moore, G. E. (1942) "A Reply to My Critics", in P. Schlipp (ed.), *The Philosophy of G. E. Moore*. Evanston: Tudor, pp. 535-677.

Neale, S. (1990) *Descriptions*. Cambridge, MA: The MIT Press.

___ (2007) "Heavy Hands, Magic, and Scene-Reading Traps", *European Journal of Analytic Philosophy* 3(2): 77-13.

Nicolle, S. and B. Clark (1999) "Experimental Pragmatics and What is Said", *Cognition* 69: 337-354.

Novacek, I. (2001) "When Children are More Logical than Adults", *Cognition* 78: 165-188.

Pritchard, D. (2010) "Contextualism, Skepticism, and Warranted Assertibility Manoeuvres," in Campbell, J.C., M. O'Rourke and H. Silverstein (eds.), *Knowledge and Skepticism*. Cambridge MA: The MIT Press, pp. 85-103.

Recanati, F. (1993) *Direct Reference*. Cambridge, MA: Blackwell.

Rysiew, P. (2001) "The Context-Sensitivity of Knowledge Attributions", *Noûs* 35(4): 477-514.

___ (2005) "Contesting Contextualism", *Grazer Philosophische Studien* 69: 51-70.

___ (2007) "Speaking of Knowing," *Noûs* 41(4): 627-662.

___ (2012) "Epistemic Scorekeeping," in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions*. Oxford: Oxford University Press, pp. 270-293.

Saul, J. (2002) "What is Said and Psychological Reality", *Linguistics and Philosophy* 25: 347-372.

Sperber, D. and D. Wilson (1981) "Pragmatics", *Cognition* 10: 281-286.

___ (1986a) *Relevance*. Cambridge, MA: Blackwell.

___ (1986b) "Loose Talk", *Proceedings of the Aristotelian Society* 6: 153-171.

___ (1987) "Précis of *Relevance*", *Behavioral and Brain Sciences* 10: 697-754.

Schaffer, J. (2004) "Skepticism, Contextualism, and Discrimination", *Philosophy and*

Phenomenological Research 69(1): 138-155.

Stanley, J. (2005) "Fallibilism and Concessive Knowledge Attributions", *Analysis* 65(2): 126-131.

Unger, P. (1975) *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.

Van der Henst, J.-B., L. Carles, and D. Sperber (2002) "Truthfulness and Relevance in Telling the Time", *Mind & Language* 17(5): 457-466.