

Factive Presupposition and the Truth Condition on Knowledge

Allan Hazlett

In “The Myth of Factive Verbs” (Hazlett 2010), I had four closely related goals. The first (pp. 497-99, p. 522) was to criticize appeals to ordinary language in epistemology. The second (p. 499) was to criticize the argument that truth is a necessary condition on knowledge because “knows” is factive. The third (pp. 507-19) – which was the intended means of achieving the first two – was to defend a semantics for “knows” on which <S knows p> can be true even if p isn’t true. The fourth (Ibid.) – which seemed necessary for the success of the third – was to defend a pragmatic account of the fact that utterances of <S knows p> typically imply p, on which the implication in those cases is down to conversational implicature.¹

In this paper I’ll go after these goals again, with an emphasis on the second. Our topic will be whether the factivity of “knows” (whatever this amounts to) supports the truth condition on knowledge. A new goal will be to defend my argument against some criticisms from John Turri (2011) and Savas Tsohatzidis (forthcoming [TO THE EDITOR: I’ve cited Tsohatzidis as “forthcoming” throughout, but I believe his paper will be appearing in the same issue of *Acta Analytica* as mine, so if there is some more elegant way of doing these citations, go ahead]). We’ll first look at the truth condition (§1) and factive presupposition (§§2 – 3), before turning to replies to Turri and Tsohatzidis (§§4 – 7).

1 The truth condition

We find the idea that knowledge is a species of true belief from the very beginning in epistemology, in Plato’s *Meno* and *Theatetus*. Why think that this? Why think, in other words, that truth is a necessary condition on knowledge?

One possibility: the truth condition is *obvious*. This is usually the defense offered in epistemology textbooks (e.g. Feldman 2003). This isn’t always obvious to people without any epistemological training, but it may well be obvious to everyone with epistemological training. But if that’s right, we might expect some account of why truth is a necessary condition on knowledge – an account of which the experts are in some way aware. Another possibility: the truth condition follows from some claim about the *state* of knowledge, for example, that knowledge is a “factive mental state,” i.e. “a propositional attitude [such that] necessarily, one has it only to truths” (Williamson 2000, p. 34, Fumerton 2006, p. 12). This premise, however, isn’t all that different from the conclusion that truth is a necessary condition on knowledge. And why think that knowledge is a factive mental state? Perhaps just because it is obvious.

A third possibility is to defend the truth condition by appeal to some claim about the *concept* of knowledge. For example, you might argue that the concept of knowledge is a “success concept” (BonJour 2002, p. 34). A fourth, related, possibility is to defend the truth condition by appeal to some claim about the *word* “knows.” For example, the claim that “knows” is factive. This is the defense of the truth condition that I want to criticize.

2 Factivity and presupposition

¹ “Implication” is used here as a neutral term, covering entailments, conversational implicatures, presuppositions, etc.

What does it mean to say that “knows” is factive? In “Myth” I understood this as the claim that an utterance of <S knows p> is true only if p. Linguists typically take the phenomenon of factivity to be a matter of *presupposition*. The notion of presupposition is controversial (cf. Karttunen and Peters 1979, Atlas 2004), but covers a wide range of cases. Consider P.F. Strawson’s (1950) take on descriptions. Two things have to happen in order for one to assert something true using a description <The F is G>: the description has to *refer* to something, the F, and what you *predicate* of the F really has to apply to the F – i.e. the F has to be G. When the first of these two things fails to happen, one’s attempt at predication fails, one’s utterance fails to express a proposition, and has no truth value (so Strawson argued). When there is no King of France, the question of whether the King of France is bald does not, and cannot, arise.

Similarly, the use of demonstrative noun phrases is said to presuppose the existence of their referents. The use of certain particles is standardly regarded as involving presupposition, for example: saying “Even Oscar enjoyed himself” presupposes that people other than Oscar enjoyed themselves. And the use of “knows” and other *factives* is standardly taken to involve presupposition.

According to a *semantic* or *lexical* account of presupposition (Frege 1892, Strawson 1950, Karttunen and Peters 1979, Katz 1979; cf. Abusch 2002), presuppositions are triggered as a matter of the conventional meaning of specific lexical items like words or sentences. Such an account of factive presupposition, applied to “knows,” would say that “knows” or <S knows p> triggers the presupposition that p is true. Recall Strawson on descriptions. We could say the same thing about factives: just as <the F> must refer for the question of the truth of <The F is G> to arise, so p must be true for the question of the truth of <S knows p> to arise. And we can take the reason to be the same, if we understand <S knows p> as having (something like) the following deep structure (Kiparsky and Kiparsky 1971, pp. 355-6; see also Glanzberg 2005, p. 369, p. 383, Holton manuscript):

S has knowledge of the fact that p.

Strawson would argue that <S> must refer, else this sentence has no truth value. On the present account, it is also the case that <the fact that p> must refer, else this sentence has no truth value. To correctly attribute knowledge of the fact that p to S, it must first be the case that both S and the fact that p exist, else <S knows p> has no truth value; once this is established, we can then answer the question of whether S and the fact that p stand in the relation of knowledge, i.e. the question of what the truth value of <S knows p> is.²

Regardless of what account of presupposition we offer, if <S knows p> presupposes p, we immediately face a puzzle, since according to a standard assumption, understood to provide one aspect of what is distinctive about presupposition, presupposition is preserved under negation. If <X> presupposes p, then <~X> presupposes p as well. “The King of France is bald” presupposes the existence of the King, but so does “The

² An alternative semantic account would be one that treats the connection between <S knows p> and p as a matter of a *conventional implicature* triggered by “knows” (cf. Karttunen and Peters 1979). The implication is not down to the truth-conditions of <S knows p>, but rather to a non-truth-conditional aspect of the conventional meaning of <S knows p>. Note that such a view doesn’t support the truth condition on knowledge.

King of France is not bald.” This seems to imply that <S knows p> doesn’t presuppose p after all, as L. Jonathan Cohen (1992) argues:

What is undeniable about an ascription of knowledge that *p* is that by uttering it the utterer also implicitly admits, affirms, or otherwise commits himself to the proposition that *p*. One cannot say – literally – ‘He knows that *p*, and it is false that *p*.’ But is this because the proposition ‘He knows that *p*’ entails the proposition that *p*? Or is it instead because the speech-act of saying ‘He knows that *p*’ normally gives the hearer to understand that the speaker is disposed to feel that *p* or has adopted the policy of premising that *p*? We can see that the latter explanation must be the correct one when we bear in mind that the verb ‘to know’ may be used factively in negative or interrogative constructions. (p. 91)

Just as <S knows p> presupposes p, <S doesn’t know p> and <Does S know p?> presuppose p. That, at least, seems sometimes to be the case:

- (1) Jack and Fred are watching football on television. Their neighbor Melissa had promised to watch the game with them, but she hasn’t arrived. Jack wonders why. Fred offers a hypothesis: “She doesn’t know that the game’s started.”³
- (2) Same situation; Jack asks: “Does Melissa know that the game’s started?”

In (1), Fred presupposes that the game has started, and in (2), Jack presupposes the same. However, so the argument goes, since neither <S doesn’t know p> nor <Does S know p?> entails p, factive presuppositions are not entailments – and so <S knows p> doesn’t entail p.

We’ll continue to assume a Strawsonian picture of presupposition on which presupposition failure yields truth-value gaps. The soundness of Cohen’s argument depends on how we conceive of entailment. On one way of conceiving of entailment, <X> entails q iff, necessarily, <X> is true only if q. Since “The King of France is Bald” has no truth value when there is no King of France, it follows, on the present conception of entailment, that “The King of France is bald” entails that the King of France exists. The same would apply to a Strawsonian account of the factivity of “knows.” It is not the case, on this conception of entailment, that presuppositions are not entailments. However, we might require something stronger for entailment. Suppose we said that <X> entails q only if, necessarily, if q is false then <X> is false. We might after all hope to be able to reject <X> as false if we discover that one of its entailments is false. Now a deep difference between presuppositions and entailments emerges, and Cohen’s argument appears to be sound.

In any event, <S doesn’t know p> can be true even if p is false, as in:

- (3) Laura is playing roulette and places an uncharacteristically large bet on 27. Shane and Conrad are observing; Shane wonders why Laura changed her strategy. Conrad: “She doesn’t know that it’s going to land on 27. How could she?”

³ Consider also a slightly less familiar use of “knows”: “She didn’t *know* that the next card would be a club,” said of someone whose true belief about the next card was not based on adequate evidence, or whatever. On such uses, see Hazlett 2009, pp. 615-19.

Intuitively, Conrad's utterance is true, even if the ball doesn't land on 27. But this case, which supports saying that <S doesn't know p> doesn't entail p, also seems to show that <S doesn't know p> doesn't *always* presuppose p. But <S knows p>, one might have thought, *always* presupposes p. What's going on here?

One possibility (cf. Frege 1919, Strawson 1952, van Fraassen 1971) is that <S doesn't know p> is ambiguous, admitting of either a wide scope or narrow scope reading, where the sentence on the wide scope reading is true iff <S knows p> is not true, and could be more perspicuously rendered as <It is not the case that S knows p>. The sentence, on the wide scope reading, does not presuppose p, and this is the right reading for Conrad's utterance in (3).

But now it seems that Cohen's argument is problematic for another reason. When <S doesn't know p> is given the wide scope reading, which is what is needed to generate a case in which the sentence is intuitively true while p is false, the sentence will no longer presuppose p. The argument was meant to drive a wedge between presupposition and entailment, by showing us a case of presupposition without entailment; it turns out that consideration of <S doesn't know p> can't do that: cases in which <S doesn't know p> intuitively doesn't entail p are also cases in which the sentence doesn't presuppose p.

The ambiguity view seems plausible at least when it comes to certain cases, for example:

- (4) Same situation as (1) and (2); Fred walks in from the kitchen and asks whether Melissa knows that the game has started. Jack: "She doesn't know that it's started, because it hasn't started – this is just a reply of last week's game."

But there is another possibility, at least for cases like (3): that even on the narrow scope reading, <S doesn't know p> doesn't always presuppose p. The sentence, we could say, has factive and non-factive *uses* – more on which below (§3).

However, consideration of <Does S know p?> can still do the work we need done: presupposition does not imply entailment. Thus the fact that an utterance of <S knows p> presupposes p doesn't imply that <S knows p> entails p (Hazlett 2010, p. 521). And this means that the factivity of "knows," understood as a matter of presupposition, doesn't suggest that truth is a necessary condition on knowledge.

Let's briefly descend from our discussion of "knows" to a discussion of knowledge. We find it plausible that truth is a necessary condition on knowledge. Is truth likewise a necessary condition on ignorance? There seems to be a sense in which it is: you cannot be ignorant without there being some fact of which you are ignorant. The right kind of connection with some fact constitutes knowledge; the lack of such connection constitutes ignorance. This idea jibes with the idea that an utterance of <S doesn't know p> presupposes p. But if truth is a necessary condition on knowledge in the same way that truth is a necessary condition on ignorance, then we should hesitate to say that truth is part of the *essence* or *nature* of knowledge, just as we would hesitate to say that truth is part of the essence or nature of ignorance. It would be better to say that truth is a necessary *precondition* of both knowledge and ignorance.

3 Pragmatic presupposition

By contrast with semantic accounts of factive presupposition, we might opt for a *pragmatic* account, on which factive presupposition is not (entirely) a matter of the

conventional meaning of, for example, “knows” or <S knows p>. Consider Robert Stalnaker’s (1999) account of presupposition, on which “it is persons rather than sentences, propositions or speech acts that have or make presuppositions.” (1999, p. 50, see also Grice 1989, Simons 2001, Glanzberg 2005) On such an account of factive presupposition, we will not say that <S knows p> presupposes p, but rather that particular utterances of <S knows p> presuppose p; more generally, we will say not that *words* are factive but that *uses* of words are factive.

On Stalnaker’s view, the pragmatic presuppositions of a speaker in a context are (roughly) the elements of the “common ground” in that context, the mutually recognized background beliefs of the speakers in that context. Given this understanding of presupposition in general, we can give the following sketch of an account of factive presupposition:

It is clear that “*x* knows that *P*” entails that *P*. It is also clear that in most cases when anyone asserts or denies that *x* knows that *P*, he presupposes that *P*. Can this latter fact be explained without building it into the semantics of the word? I think it can. Suppose a speaker were to assert that *x* knows that *P* in a context where the truth of *P* is in doubt or dispute. [...] He would be leaving it unclear whether his main point was to make the claim about the truth of *P*, or to make a claim about the epistemic situation of *x* (the knower), and thus leaving unclear what direction he intended or expected the conversation to take. Thus, given what “*x* knows that *P*” means ... it would be unreasonable to assert that *x* knows that *P* in such a context. (1999 p. 55)

The same, *mutatis mutandis*, for <S doesn’t know p>. This account, along with other pragmatic accounts (Grice 1989, p. 279, Atlas 2004, p. 34, Glanzberg 2005, p. 382), assumes that <S knows p> entails p, and indeed uses this assumption in giving an account of the fact that someone who says <S knows p> or <S doesn’t know p>, at least typically, presupposes p.

The appeal of pragmatic accounts comes from the fact that the behavior of paradigm factives (“knows,” “realizes,” “proves,” “discovers,” “learns,” “shows”) varies substantially with context.⁴ This is most obvious in negative, interrogative, and modal constructions (Karttunen 1971, Simons 2001, Abusch 2002, Abbott 2006). Someone’s utterance of the same sentence can involve a factive presupposition in one context but not in another. Consider negative uses of “discovers”:

- (5) Of an explorer who “went off to someplace expecting to discover that the natives were very interesting in certain respects” (Grice 1989, p. 279), but where it turned out there were no interesting natives: “He did not discover that the natives were interesting, because there were no interesting natives there.”
- (6) Of an explorer who was too culturally insensitive to appreciate how interesting the natives actually were: “He did not discover that the natives were interesting, because he did not understand the significance of their headdresses.”

⁴ Cf. The need for a division of factives, e.g. into those with soft triggers and those with hard triggers (Abusch 2002, cf. Abbott 2006) or into thin factives and thick factives (Holton ms).

In (5) the speaker does not presuppose that the natives were interesting, but in (6) the speaker does presuppose that the natives were interesting. This motivates the pragmatic view, on which some uses of “discover” are factive, and others aren’t. On my view (see §4), the same is true of “knows.”

It is important for our purposes here to note that some paradigm factives have non-factive *positive* uses:

- (7) Evidence that ulcers are caused by bacterial infection didn’t exist until the 1980s; before then, people believed ulcers were caused by stress. A doctor to her patient, in 1975: “In medical school, I learned that stress causes ulcers, so you’ll need to avoid faculty meetings for a while.”
- (8) A doctor to her colleague, in 1995: “In medical school, I learned that stress causes ulcers. Can you believe the crap they were teaching us back then?”

In (7) the doctor presupposes that stress causes ulcers, but in (8) she doesn’t. It is of course possible to argue that the doctor in (8) is not using “learns” literally, that she did not really *learn* that stress causes ulcers, because one cannot *learn* something false. Perhaps she engages in “protagonist projection” (Holton 1997, p. 626), more on which below (§4). But the appropriateness of the doctor’s utterance in (8) is defeasible, *prima facie* evidence that the factive presuppositions involved with “learns” are pragmatic, not semantic.

A final point in favor of a pragmatic approach. There seem to be both factive and non-factive uses of words that no one would call “factive verbs,” e.g. “tells”:

- (9) Dugald, David, and Duncan are tunneling, from the south, into the vault of the Aberfeldy Municipal Credit Union. When Duncan is arrested and interrogated, he cracks and lets slip that there’s a tunnel, but doesn’t reveal the location. Dugald to David: “He didn’t tell them that the tunnel’s on the south side, so we’re still in business.”
- (10) Dugald, David, and Duncan are tunneling, from the north, into the vault of the Aberfeldy Municipal Credit Union. When Duncan gets arrested on an unrelated matter, the three hoods conspire to mislead the police about the location of the tunnel; Duncan is to tell investigators that the tunnel is on the south side of the building. But Dugald finds out that Duncan cracked under pressure, and says to David: “He didn’t tell them that the tunnel’s on the south side, so we’re in trouble.”

In (9), Dugald presupposes that the tunnel is on the south side, but in (10), he doesn’t. This further suggests that the phenomenon of factive presupposition is a pragmatic matter; it is not down, in these cases, to the semantics of “tells.”

None of this is meant to be decisive in favor of a pragmatic account, as against a semantic account. I have sketched very simple versions of both semantic and pragmatic approaches (for more sophisticated versions, see Simons 2001, Abusch 2002, Atlas 2004, Glanzberg 2005, Abbott 2006). We’ll now turn to the particular pragmatic approach that I proposed in “Myth,” and to the question of whether <S knows p> entails p.

4 Cancellation and entailment

In “Myth” I argued that when an utterance of <S knows p> implies p, it’s down to conversational implicature: the speaker presupposes p in (something like) Stalnaker’s pragmatic sense. Turri (2011, p. 11) and Tsohatzidis (forthcoming, pp. 13-2) criticize this by arguing that the implication in question fails the “cancelation test” for conversational implicature. What is the test, and do the implications in question fail it?

Grice (1989) characterizes the test as follows:

[A] putative conversational implicature is explicitly cancelable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add *but not p*, or *I do not mean to imply that p*, and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. [A]ll conversational implicatures are cancelable. (p. 44)

And so if a putative conversational implicature (hereafter, “implicature”) is not cancelable, then it is not actually an implicature.

Note that what is required is that, when the implicature is canceled, the utterance is still linguistically acceptable. What is needed is that the form of words can be acceptably used without implying the putative *implicatum*. But the test is still ambiguous. Suppose someone maintains that by uttering <S> in context C, A implicates q. Two questions immediately arise:

- (i) When we test for cancelability, must we consider the specific sentence that A uttered, or can we consider sentences merely having, in some sense, the same form?
- (ii) When we test for cancelability, must we consider only C, or can we consider other contexts as well?

With contextual cancelation, of course, the answer to (ii) is obvious: a putative implicature is contextually cancelable if we can find *some* context in which the utterance of the form of words would simply not carry the implication. But with explicit cancelation, it isn’t. Consider:

- (11) On the subway, Jamie asks Anne if the train they’re on stops at 18th street. Anne can’t remember whether there even *is* an 18th street stop on this line, but she knows it’s a local train, so she says: “I know it’s a local train.”

On my view, Anne implicates that it is, indeed, a local train. Is this putative implicature explicitly cancelable?

On the one hand, you might answer this by considering a situation like (11) in all respects, but where Anne says “I know it’s a local train, but it isn’t,” and ask whether her utterance in that case would be acceptable and not imply that it’s a local train. However, Matt Weiner (2006) has shown that this kind of cancelability is not necessary for genuine implicature. Weiner’s example:

Alice and Sarah are in a crowded train; Alice ... is sprawled across two seats, and Sarah is standing. Sarah says to Alice, ‘I’m curious as to whether it would be physically possible for you to make room for someone else to sit down’. (p. 128)

Sarah has implicated (rudely) that Alice should make room for her (Sarah) to sit down. But even if Sarah were to add “Not that you should make room; I’m just curious,” she would still imply that Alice should make room. This seems like an obvious case of implicature – if there is such a thing, this is a case of it – but the implicature cannot be explicitly canceled in the original context of utterance.

We should conclude, then, that the *most* that should be required, for it to be the case that by uttering <S> in context C, A implicates q, is that there be some context C* in which an utterance of <S> would be acceptable and not imply q. Anne, in C, could not appropriately utter “I know it’s a local train, but it isn’t” – this would strike us as unacceptable. But this doesn’t show that she doesn’t implicate that it’s a local train.

But what about question (i)? When we consider whether Anne’s putative implicature in (11) is cancelable, must we search for contexts in which someone acceptably utters “I know it’s a local train” without implying that it’s a local train, or where someone acceptably utters “I know it’s a local train, but it isn’t,” or must we merely seek contexts in which someone acceptably utters <S knows p> (i.e. a sentence of the same *form* as the original) without implying p, or where someone acceptably utters <S knows p, but not-p> (i.e. explicit cancellation incorporating a sentence of the same form as the original)?

There are contexts in which an utterance of “I’m curious as to whether it would be physically possible for you to make room for someone else to sit down” doesn’t implicate that the hearer should make room – as Weiner (2006) suggests, “a madwoman or a philosopher concerned with free will” could “really be curious about whether it was physically possible for Alice to make room.” (p. 129) Are there likewise contexts in which someone could acceptably say “I know it’s a local train” without implying that it’s a local train, or where someone could acceptably say “I know it’s a local train, but it isn’t”? Such situations are hard to imagine.

However, this can be explained without conceding that the typical implication from <S knows p> to p is not down to implicature. This is because we can explain why it is never acceptable to say “I know it’s a local train, but it isn’t” without appealing to the factivity of “knows.” Belief is a necessary condition on knowledge, and it is an obvious necessary condition, such that anyone who asserts that S knows p is openly committed to S’s believing p. But if so, then an utterance of <I know p, but not-p> will always strike us as inappropriate, because the speaker has committed herself to believing p, while in the same breath committed herself to p being false. Such an utterance will always have a “paradoxical” aspect, akin to that enjoyed by Moore’s famous sentence, “Though I don’t believe it’s raining, yet as a matter of fact it really is raining.” (1993, p. 207)

For this reason, we should consider third-person knowledge attributions, or first-person past-tense knowledge attributions. Consider:

- (12) Said as part of lecture on the history of ulcer prevention techniques, in 1975, discussing the techniques of the past: “Everybody knew that stress causes ulcers; what they didn’t now was how to reduce stress.”

The speaker in (12), on my view, implicates that stress causes ulcers. This proposition is part of the common ground, given what she and her audience mutually recognize one another as believing; in the pragmatic sense (§3) the speaker presupposes that stress causes ulcers. However, there are contexts in which the sentence “Everybody knew that

stress causes ulcers” can be acceptably uttered and where it doesn’t imply that stress causes ulcers. Consider:

- (13) Same set up as (12), now said as part of a lecture on the history of ulcer prevention techniques, in 1995, discussing changing knowledge of ulcers:
“Everybody knew that stress causes ulcers, before more advanced work was done in the 1980s. They’re actually caused by bacterial infection.”⁵

So the putative implicature in (12), at least, is explicitly cancelable, in the relevant sense.

Now you might wonder whether the test even requires this much. For Grice speaks of finding a context in which an utterance of the “form of words” doesn’t carry the implicature. So perhaps all we need is to see if utterances of <S knows p> can ever be acceptable and not imply p, or whether utterances of <S knows p, but not-p> can ever be acceptable.

Indeed, this will get us to the heart of the matter: to find out whether <S knows p> entails p, our best bet is to see if an utterance of <S knows p, but not-p> can ever be acceptable. But this is what non-factive uses of “knows” provide. Almost everyone who reflects on ordinary usage agrees that sometimes it is acceptable to utter <S knows p, but not-p>. Such cases are so familiar that many introductory epistemology textbooks talk about them (e.g. BonJour 2002, p. 33, Feldman 2003, p. 13). But it only takes one case, like (13), to see the point.⁶ This is defeasible, prima facie evidence that <S knows p> doesn’t entail p.

Now it is possible that there are no cases that will convince you of this. It isn’t hard to explain cases away on a case-by-case basis. As Turri argues (forthcoming, p. 6), the acceptability of “I knew q was going to happen,” in contexts in which it is clear that q did not happen, are not cases in which “knows” is used non-factively, because that q was going to happen and that q did not happen are compatible – it was *true* that q was going to happen, but then it didn’t.⁷ Those who say “People knew the earth was flat” are often inclined to say that it was *true* that the earth was flat, so this kind of case isn’t obviously a non-factive use either.

I suspect that the genuine cases in which “knows” is used non-factively are cases in which we mean to indicate that something false was *obvious*. We are particularly inclined to use “knows” non-factively when we are speaking of conventional wisdom or common knowledge – i.e. when we are using “knowledge” in the way that sociologists use it, to

⁵ Cf. “Myth,” p. 501; the sentence is adapted from one appearing in “Cat Carrier: Your Cat Could Make You Crazy,” *National Geographic* 208, July 2005.

⁶ Objection: in (11) the doctor says <S knew p, but not-p>, not <S knows p, but not-p>. Reply: if <S knows p, but not-p> entails p, but <S knew p, but not-p> doesn’t entail p, then there is no support here for the view that truth is a necessary condition on knowledge. What this would show is that the factive presuppositions involved with utterances of <S knows p> are not down to something about the nature of knowledge.

⁷ This point is worth conceding, but it’s not obviously right. Corporal Dakota Meyer is said to have described a narrow escape in Afghanistan by saying, “I didn’t think I was going to die. I knew I was” (“Top Medal for Marine who Saved Many Lives,” *New York Times*, 15th September, 2011). The details of Corporal Meyer’s conversation with the President aren’t available, but it’s easy to imagine the context such that Meyer is suggesting that his knowledge turned out to be false.

describe the knowledge of a community, culture, or institution (cf. Goldman 2002, pp. 183-5, and Kusch 2009, pp. 72-3). This is also the way we speak of scientific knowledge, as when we speak of the progress of scientific knowledge over time, or when we speak of our knowledge improving. Consider:

- (14) From “Aid 2.0,” *The Economist*, 13th August, 2011: “Every so often something comes along which shows that almost everything you know about a subject is wrong.”⁸

The utterance in (13) is also one of this kind. Many of these cases will be cases of obviousness within a community, but there are cases where the obviousness in question is relative to an individual. Richard Holton (1997, p. 626) cites this passage:

- (15) From Ian McEwan, *The Innocent*: “He knew from experience that unless he made a formidable effort, a pattern was waiting to impose itself: a polite enquiry would elicit a polite response and another question ... He had asked her about tea making. One more like that, and there would be nothing he could do ... Rather than tolerate more silence he settled after all for small talk, and began to ask, ‘Have you lived here long?’ But all in a rush she spoke over him, saying, ‘How do you look without your glasses? Show me please.’”

Here’s my take on all this: there are factive and non-factive uses of “knows.” Sometimes – indeed most of the time – an utterance of <S knows p> implicates p (“Myth,” pp. 511-15); the speaker in such cases presupposes p. Likewise, much of the time, though perhaps less of the time, an utterance of <S doesn’t know p> implicates p. When an utterance of <S knows p> does implicate p, there is no reason to expect that said utterance could be replaced, in the same context, with an utterance of <S knows p, but not-p>; that was the point from Weiner (2006). There is no reason to think that implicatures are cancelable in *that* sense. They must be cancelable in a weaker sense, however: if I say that utterances of <S knows p> typically implicate p, then it must be possible to find cases in which <S knows p> is acceptable and doesn’t imply p, or it must be possible to find cases in which <S knows p, but not-p> is acceptable. But it seems to me that we can find such cases. This is defeasible, prima facie evidence that <S knows p> doesn’t entail p.

Suppose I am right that <S knows p> doesn’t entail p, and that there are factive and non-factive uses of “knows.” If so, we need not worry about Tsohatzidis’ objection (forthcoming, p. 6 and passim) that some putative examples of non-factive uses of “knows,” such as (12), require other words (in this case, “shows”) to be factive – in (12), “knows” is used non-factively and “shows” is used factively. It is uses of words, not words, that are factive, on the pragmatic view I favor; speakers, and not words, make factive presuppositions.

An important objection (DeRose 2009, pp. 15-17, Hannon forthcoming) to the claim that there are non-factive positive uses of “knows” is the view that all would-be non-factive positive uses of “knows” are cases of what Richard Holton (1997) calls “protagonist projection.” (p. 626) If A says to B, “He gave her a ring made of diamonds, but the diamonds turned out to be made of glass,” both parties understand not that there

⁸ Turri (2011, p. 6) suggests that “wrong” doesn’t mean false in these contexts. I find this unpersuasive: the relevant utterances would sound just as acceptable with “false” in place of “wrong.”

can be glass diamonds but that the stones seemed to be diamonds, from the perspective of the protagonist of the story. So the argument goes, if A says to B, <S knows p, but not-p>, both parties understand not that there can be false knowledge but that p seemed to be true, from the perspective of the protagonist of the story. It is essential that in all cases of protagonist projection, both speaker and hearer understand that the relevant utterance is false. So, for example, when A says to B, “The diamonds turned out to be made of glass,” for this to be a case of protagonist projection both A and B must know that diamonds can’t be made of glass. For this reason, when A says to B, “The Ancients knew that the earth was flat,” for this to be a case of protagonist projection both A and B must know that false propositions can’t be known. For my part, it is unclear whether ordinary speakers of English generally know that false propositions can’t be known – which explains their somewhat muddled responses to the truth condition in our introductory epistemology courses. For this reason, it is unclear whether all would-be non-factive positive uses of “knows” are cases of protagonist projection.

Wesley Buckwalter (ms) devised an experiment in which subjects were asked, of the sentence “Everyone knew that stress caused ulcers, before two Australian doctors in the early 1980s proved that ulcers are actually caused by bacterial infection,” whether “everyone thought they knew” or whether “everyone really knew,” and found that the great majority of subjects said that “everyone thought they knew,” rather than saying that “everyone really knew.” I don’t think this is good evidence that all would-be non-factive positive uses of “knows” are cases of protagonist projection (and Buckwalter doesn’t claim that it is), although further inquiry is surely called for. The subjects obviously made the correct choice in answering as they did. One reason for this might have been that the most natural reading of the question is as a question about whether stress causes ulcers: to answer that “everyone really knew” would imply, given the form of the question, that stress causes ulcers, while to answer that “everyone thought they knew” would make clear that the subject had understood that stress doesn’t cause ulcers (for more on <S thinks she knows p>, see §5).

Propositional attitude attributions can be used to do either of two rather different things. On the one hand, they can be used to describe someone’s state of mind. On the other, they can be used to imply information about the world. In some contexts, conversational participants are concerned only with people’s states of mind. In others, they inquire about the world, e.g. about whether p is true, and in such contexts the attribution of knowledge will implicate that the known proposition is true. Consider a word like “learns,” and a non-factive use of learns: “In school we learned that World War I was a war to ‘make the world safe for democracy,’ when it was really a war to make the world safe for the Western imperial powers.” It seems to me that the question of whether “they thought they learned” or whether “they really did learn” is ambiguous: it could be a question about the students’ states of mind (are the students misremembering their curriculum, or is that really what they were taught?) or about the world (was their lesson just pernicious and false ideology, or did World War I actually make the world safe for democracy?). I speculate that this ambiguity is partly behind Buckwalter’s results: when subjects were asked this question: about half said that “they thought they learned” and the other half said that “they really learned” (Ibid.).⁹

Consider the claim that the Ancients knew that the earth was flat. It seems to me that we might ask whether the Ancients *really* knew that the earth was flat with two different

⁹ Still needed, however, is an explanation for the difference in their responses to the “learns” item and the “knows” item.

contrasts in mind, which might be made explicit. We might ask whether the Ancients really knew that the earth was flat, as opposed to the earth being round. And here the obviously correct answer is that the earth is round, and that they did not really know. But we also might ask whether the Ancients really knew that the earth was flat, as opposed to this not being an item of common knowledge in Ancient times, or as opposed to this being more of a hypothesis than a settled conviction, and here the answer isn't obvious, because it's a question about the socio-anthropology of the relevant cultures, and not a question about whether the earth is round.

Consider a familiar case of conversational implicature: the local says that there's a gas station around the corner, but when you arrive it turns out that its shut down all month for repairs. The local implied that the gas station was open, but this was not entailed by what he said. However, take the question "Was there *really* was a gas station around the corner, or did the local mislead you?" The obviously correct answer to *this* question is that the local misled you, rather than saying that there really was a gas station around the corner. This doesn't mean that it's not the case that there was a gas station around the corner. Likewise, I am proposing, the fact that Buckwalter's subjects said that "everyone thought they knew," rather than saying that "everyone really know," doesn't mean that it's not the case that everyone knew.

5 Two questions

Tsohatzidis writes:

[S]omeone who ... is willing to regard the acceptability of certain utterances of the form "S knows that p, but \sim p" as evidence that what those utterances express is not contradictory, would be expected to be equally willing to regard the *un*acceptability of certain other utterances of the form "S knows that p, but \sim p" as evidence that what these *other* utterances express *is* contradictory. (forthcoming, p. 7)

It seems to me (cf. "Myth," p. 507) that the fact that utterances of <S knows p, but not-p> are sometimes acceptable is defeasible, prima facie evidence that <S knows p, but not-p> is not contradictory. This is because, on my view, the fact an utterance is acceptable is defeasible, prima facie evidence that it is true, and that therefore the sentence uttered is not a contradiction. This does not, it seems to me, suggest the move that Tsohatzidis suggests: that of taking the fact that an utterance is unacceptable as defeasible, prima facie evidence that the sentence uttered is a contradiction. For an utterance can be unacceptable for a plurality of reasons, most of which do not involve the sentence uttered being contradictory: it may be contingently false, or rude, or irrelevant, and so on.

However, sometimes – maybe most of the time – an utterance of <S knows p, but not-p> is unacceptable. We could explain this by appeal to the fact that <S knows p> entails p. If that's not the case, why is <S knows p, but not-p> unacceptable, when it's unacceptable? This is our first question (cf. Turri 2011, p. 8).

We have already given an account of this in the case of first-person present-tense knowledge attributions. The account I favor in other cases appeals to the norms of cooperation proposed by Grice (1989), and in particular the requirement of relevance. Consider:

- (16) Detectives are investigating a bombing. A asks B whether the FBI have provided any information about the materials used to make the bomb. B replies: “They know that the bomb was homemade.” (“Myth,” p. 510)

Here, I argued, B uses “knows” to “relay” the testimony of the FBI to A. But now consider:

- (17) Same as (16). The bomb was not, however, homemade, and B knows this. But it is obvious to the FBI that it was homemade – obvious in whatever way it was obvious to people that stress causes ulcers. In response to A’s question, B says: “They know that the bomb was homemade, but it wasn’t.”

B’s utterance here strikes us as unacceptable. The reason, I propose, is that it violates the requirement of relevance. What A wants to know about is the composition of the bomb, not the FBI’s state of mind. For B to use “knows” non-factively here is to offer A irrelevant information – irrelevant and, if not for A’s final point, potentially misleading information.

Recall the idea (§4) that propositional attitude ascriptions can be used either to talk about someone’s state of mind or to talk about the world. In the latter sort of context, the *non-factive* attribution of knowledge that p will be irrelevant, and therefore unacceptable. This irrelevance is not mitigated by the attempted cancellation, any more than saying “I know I’m being rude” mitigates one’s rudeness.

Something similar can be said, *mutatis mutandis*, of contexts in which p is something all parties assume, and their inquiry concerns people’s states of mind vis-à-vis p. If one of the parties to such a conversation has come to reject p, then saying <S knows p, but not-p> is uncooperative. The attribution of knowledge to S is inappropriate here, given the speaker’s change of heart about p. <S knows p, but not-p> is not a perspicuous way to tell your conversational partners that you have come to disagree with them about p. The mention of S’s mental state is inappropriate, given the speaker’s true conversational purpose, and so her utterance strikes us as unacceptable.

This argument assumes that our intuitions about unacceptability are neutral as to the species of unacceptability (cf. “Myth,” p. 498n). Our linguistic intuitions come in the form *that utterance is unacceptable*, not in the form *that sentence is contradictory*. Just how unacceptable an irrelevant utterance seems will come down the degree to which the speaker is being uncooperative. In (17), I propose, B is being highly uncooperative.

Let’s turn now to our second question. Someone who utters <S thinks that she knows p> often implies that p is false (cf. Turri 2011, p. 7, Hannon forthcoming). We could explain this by appeal to the fact that <S knows p> entails p: if p is false, then S doesn’t know p, and thus must merely think that she knows p. If <S knows p> doesn’t entail p, why does an utterance of utters <S thinks that she knows p> typically imply that p is false?

Our answer again appeals to Gricean norms of cooperation. Consider a context in which the participants are in doubt as to whether p is true or false, and the question has arisen as to whether S knows whether p. Saying <S thinks that she knows p>, in such a context, would implicate that p is false, because the other participants would need to figure out why the speaker has brought up S’s meta-beliefs about whether she knows p, rather than just talking about whether S knows p. The best explanation of this is that the

speaker thinks that *p* is false. But in other contexts where the question of whether *S* knows *p* is salient, such as a context in which the participants are assuming that *p* is true, an utterance of <*S* thinks that she knows *p*> will not imply that *p* is false. We should keep in mind that, on the view I'm proposing, "knows" has factive and non-factive uses – utterances of <*S* thinks that she knows *p*> that imply that *p* is false may be understood as cases in which "knows" is used factively.

6 Stalnaker again

There is a puzzle for Stalnaker's view of factive presupposition (§3). Stalnaker, and other defenders of pragmatic accounts, maintain that someone who says <*S* knows *p*> implicates *p*. But they also maintain that <*S* knows *p*> entails *p*. However, if <*S* knows *p*> entails *p*, then it is never acceptable to utter <*S* knows *p*, but not-*p*>. But if that is never acceptable, then it seems that their purported implicature fails the cancelation test, and is not a genuine implicature.

If so, as Jay David Atlas points out (2004, pp. 38-9), Stalnaker's pragmatic account of factive presupposition just sketched diverges importantly from semantic accounts. For a semantic account, there is no particular asymmetry between the presuppositions involved in uttering <*S* knows *p*> and <*S* doesn't know *p*> (on the narrow scope reading). But on the pragmatic account just sketched, there is a deep asymmetry here: when someone presupposes that *p* by uttering <*S* doesn't know *p*>, it is down to implicature, but when someone presupposes that *p* by uttering <*S* knows *p*> it cannot be down to implicature – because the implicature cannot be canceled, given that <*S* knows *p*> entails *p*.

The account I favor gives a unified explanation of the presuppositions typically involved in uttering <*S* knows *p*> and <*S* doesn't know *p*> (on the narrow scope reading), since both presuppositions are down to implicature. Unified explanations are appealing. I think this is what Cohen was getting at (§2): he is assuming that factive presuppositions are all members of some common kind, and so if some of them don't involve entailment (as with utterances of <*S* doesn't know *p*>, then the others (e.g. those involved in utterances of <*S* knows *p*>) don't involve entailment, either.

This seems to present a dilemma for the defender of a pragmatic account that would resist my argument that <*S* knows *p*> doesn't entail *p*. If the defender of the pragmatic account says that <*S* knows *p*> entails *p*, then her purported implicatures fail the cancelation test. On the other hand, if she were to reject the cancelation test, then the principal objection to my argument is undermined.

7 "Knows" as polysemous

Tsohatzidis (forthcoming) argues "that "knows" has two distinct senses, a factive and a non-factive one." (p. 15) Turri (2011) suggests something similar, suggesting a distinction between the concept of "knowing" and the concept of "knowing something is true." (p. 14) Let us suppose that this is right: "knows" is polysemous, having both a factive and a non-factive sense. (This is the semantic analogue of my pragmatic view that "knows" has factive and non-factive uses, §4)

Recall my second goal in "Myth": to criticize the argument that truth is a necessary condition on knowledge because "knows" is factive. If "knows" is polysemous in the way that Tsohatzidis claims, then the factivity of "knows" does not support the truth condition. As he argues:

[F]active knowledge is the chosen subject matter of epistemology, and the fact that, besides its factive sense, the word “know” has another sense which is not factive neither augments nor diminishes the linguistic plausibility of the traditional epistemologist’s choice of subject matter: it is simply linguistically irrelevant to the choice. (p. 20)

For the same reason, the fact that the word “knows” has a factive sense is linguistically irrelevant to the traditional epistemologist’s choice of subject matter.

Now in “Myth” (p. 500) I wanted to conclude from this that epistemological claims *never* have “linguistic plausibility” – I proposed a “divorce” between the epistemological theory of knowledge and the linguistic theory of knowledge attributions. That was my first goal in “Myth”: to criticize appeals to ordinary language in epistemology. Turri (2011) argues that this conclusion was too quick, for if the distinction between “knowing” and “knowing something is true” is legitimate, then:

... we may charitably understand philosophers to be interested in the concept expressed by ‘know it is true’, and epistemologists may properly appeal to usage of ‘knows it is true’ and the concept it ordinarily expresses to constrain and guide epistemological theorizing. (p. 14)

The idea here is not to appeal to ordinary language to motivate the truth condition, but to motivate other epistemological choices by appeal to ordinary language. This will work, so long as we can distinguish ordinary usage of the concept expressed by “knows it is true” and the concept expressed by (non-factive uses of) “knows.” But this is not a problem so long as it’s relatively easy to tell which sense of “knows” is being used in any given situation. Whether ordinary language appeals in epistemology are legitimate will depend on how the practice is conducted – it will require a lot of caution and care to make sure that only the relevant sense of “knows” is being employed. But Turri’s point is apt: there aren’t grounds for divorce, at least not yet.

Perhaps, however, a trial separation. The appeal of a non-linguistic approach can be bolstered by considering the question of how the truth condition might, after all, be motivated, if not on linguistic grounds. One possibility is that it is just a matter of epistemological convention, or perhaps of something a bit less arbitrary: the interests of philosophers, as Turri (2011, p. 14) and Tsohatzidis (forthcoming, p. 21) suggest. Epistemologists often set aside (apparently) non-propositional species of knowledge – acquaintance knowledge, practical knowledge – in favor of propositional knowledge, so perhaps non-factive propositional knowledge should simply be added to this list of species to be set aside.

But why? Tsohatzidis imagines a field of study devoted to non-factive propositional knowledge, but “cannot see what the interest of establishing such a field might be.” (p. 21) As Martin Kusch (2009) argues, “historians and sociologists of knowledge” use “knowledge” in a non-factive way, in their attempt to “develop theories that explain why certain types of belief – or indeed, certain types of *knowledge* – are found credible with certain types of communities.” (p. 73) For my part, that such theorists are interested in such inquires seems sufficient to absolve them of criticism; I can’t imagine how I could justify *my* interest in the things that *I* study. And it would surely be a mistake (not one that Tsohatzidis suggests, of course) to accuse these theorists of misspeaking when they use “knowledge” in this way. It would be “a misplaced attempt at linguistic legislating.” (Ibid.)

A final possibility. Might factive (propositional) knowledge be more valuable than non-factive (propositional) knowledge? Perhaps the choice to study factive knowledge is justified in light of the value of factive knowledge, above that of non-factive knowledge, just as (so you might argue) the choice to study knowledge is justified in light of the value of knowledge, above that of true belief.

We shall have to tread carefully here. When we set aside acquaintance knowledge and practical knowledge, there is no implication that these are less valuable than propositional knowledge. To choose to study x might commit one to the value of knowledge about x , but it does not commit one to the value of x (consider the scatologist), although some epistemological fans of understanding have missed this. The value of x is not necessary to justify the study of x . But might it be sufficient?

The most salient difference between factive knowledge and non-factive knowledge, of course, is that factive knowledge is a species of true belief – thus suggesting that the presence of true belief somehow makes for the difference in value between factive and non-factive knowledge. Knowledge, some argue (Sosa 2007, Greco 2010), is a valuable cognitive achievement, the acquisition of truth through intellectual virtue, valuable both in virtue of the value of the object acquired, viz. truth, and in virtue of the value of achievement in general. Non-factive knowledge is no such valuable achievement, as truth is not acquired. This might justify our study of factive knowledge, as against non-factive knowledge.

But where do we now stand vis-à-vis the methodology of appealing to ordinary language in epistemology? We have arrived at a conception of knowledge as a valuable cognitive achievement by appeal to the value of truth and the value of achievement. What role should, or could, ordinary language play in this inquiry? Is our conclusion threatened if it fails to correspond to some sense of “knows” in ordinary language? No, for that would not undermine our evaluative premises, about truth and achievement. Is our conclusion supported if it does so correspond? Perhaps in some sense; if some sense of “knows” in ordinary language corresponds to our conception of knowledge as a valuable cognitive achievement, our conception might be further refined upon examination of that sense. It seems to me to be a very open question whether there is any sense of “knows” in ordinary language that corresponds to this conception. In any event, once such evaluative considerations are given center stage, the epistemological importance of that question is diminished.¹⁰

Bibliography

Abbott, B. (2006), “Where Have Some of the Presuppositions Gone?,” in B. Birner and G. Ward (eds.), *Drawing the Boundaries of Meaning: Neo-Gricean Studies in Pragmatics and Semantics in Honor of Laurence R. Horn* (John Benjamins), pp. 1-30.

Abusch, D. (2002), “Lexical Alternatives as a Source of Pragmatic Presuppositions,” in B. Jackson (ed.), *Proceedings of SALT XII* (CLC Publications), pp. 1-20.

Atlas, J.D. (2004), “Presupposition,” in L.R. Horn and G. Ward (eds.), *The Handbook of Pragmatics* (Blackwell), pp. 29-52.

¹⁰ Thanks to Wesley Buckwalter, Matthew Chrisman, Richard Holton, Jennifer Nagel, and Duncan Pritchard for valuable comments and discussions.

- BonJour, L. (2002), *Epistemology: Classic Problems and Contemporary Responses* (Rowman & Littlefield).
- Buckwalter, W. (ms), “Factive Verbs and Protagonist Projection.”
- Cohen, L.J. (1992), *An Essay on Belief and Acceptance* (Oxford University Press).
- Feldman, R. (2003), *Epistemology* (Prentice-Hall).
- Frege, G. (1892), “Über Sinn und Bedeutung,” translated as “On Sense and Reference” in *Translations from the Philosophical Writings of Gottlob Frege* (Oxford University Press, 1970), pp. 56-78.
- (1919), “Negation,” reprinted in *Translations from the Philosophical Writings of Gottlob Frege* (Oxford University Press, 1970), pp. 117-35.
- Fumerton, R. (2006), *Epistemology* (Blackwell).
- Glanzberg, M. (2005), “Presupposition, Truth Values, and Expressing Propositions,” in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Truth, and Meaning* (Oxford University Press), pp. 349-96.
- Greco, J. (2010), *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity* (Oxford University Press).
- Grice, H.P. (1989) *Studies in the Way of Words* (Harvard University Press).
- Hannon, M. (forthcoming), “‘Knows’ Entails ‘Truth,’” *Journal of Philosophical Research*.
- Hazlett, A. (2009), “Knowledge and Conversation,” *Philosophy and Phenomenological Research* 78:3, pp. 591-620.
- (2010), “The Myth of Factive Verbs,” *Philosophy and Phenomenological Research* 80:3, pp. 497-522.
- Holton, R. (1997), “Some Telling Examples: A Reply to Tsohatzidis,” *Journal of Pragmatics* 28, pp. 625-28.
- (manuscript), “Facts, Factives, and Contra-Factives.”
- Karttunen, L. (1971), “Some Observations on Factivity,” *Papers in Linguistics* 5, pp. 55-69.
- Karttunen, L., and Peters, L. (1979), “Conventional Implicature,” in C. Oh and D.A. Dinneen (eds.), *Syntax and Semantics, Volume 11: Presupposition* (Academic Press), pp. 1-56.
- Katz, J. (1979), “A Solution to the Projection Problem for Presupposition,” in C. Oh and D.A. Dinneen (eds.), *Syntax and Semantics, Volume 11: Presupposition* (Academic Press), pp. 91-126
- Kiparsky, P., and Kiparsky, C. (1971), “Fact,” in D.D. Steinberg and L.A. Jakobovits (eds.), *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology* (Cambridge University Press), pp. 345-69.

- Kusch, M. (2009), "Testimony and the Value of Knowledge," in A. Haddock, A. Milar, and D. Pitchard (eds.), *Epistemic Value* (Oxford University Press), pp. 60-94.
- Simons, M. (2001), "On the Conversational Basis of Some Presuppositions," in R. Hastings, B. Jackson, and Z. Zvolensky (eds.), *Proceedings of Semantics and Linguistics Theory 11* (CLC Publications), pp. 431-48.
- Sosa, E. (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (Oxford University Press).
- Stalnaker, R. (1999), "Pragmatic Presuppositions," in his *Context and Content* (Oxford University Press), pp. 45-62.
- Strawson, P.F. (1950), "On Referring," *Mind* 59, pp. 230-44.
- (1952), *Introduction to Logical Theory* (Methuen).
- Tsohatzidis, S. (forthcoming), "How to Forget that "Know" is Factive," *Acta Analytica*.
- Turri, J. (2011), "Mythology of the Factive," *Logos & Episteme* 2:1, pp. 143-52.
- van Fraassen, B.C. (1971), *Formal Semantics and Logic* (Macmillan).
- Williamson, T. (2000), *Knowledge and Its Limits* (Oxford University Press).
- Weiner, M. (2006), "Are All Conversational Implicatures Cancellable?," *Analysis* 66:2, pp. 127-30.