

'Knows' Entails Truth

Michael Hannon
University of Cambridge

King's College, Cambridge
CB2 1ST, United Kingdom

mh536@cam.ac.uk

Abstract

It is almost universally presumed that knowledge is *factive*: in order to know that p it must be the case that p is true. This idea is often justified by appealing to knowledge ascriptions and related linguistic phenomena; i.e., an utterance of the form 'S knows that p , but not- p ' sounds contradictory. In a recent article, Allan Hazlett argues that our ordinary concept of knowledge is *not* factive. From this it seems to follow that epistemologists cannot appeal to ordinary language to justify the truth condition of knowledge. More significantly, Hazlett claims that epistemologists theorizing about knowledge should not concern themselves with the ordinary concept of knowledge as revealed by knowledge ascriptions and related linguistic phenomena. My paper has two goals: first, to defend the orthodox view that the ordinary concept of knowledge is factive; second, to undermine Hazlett's claim that epistemologists should not theorize about knowledge on the basis of how 'knows' is used in everyday speech.

I. INTRODUCTION

According to what is often referred to as the ‘traditional’ analysis of knowledge, a subject S knows some proposition p if and only if

- (i) S believes that p ;
- (ii) S is sufficiently justified in believing that p ;
- (iii) p is true.

Few people currently accept these conditions as sufficient for knowledge (see Gettier 1963). There has also been controversy over whether knowledge requires belief (Radford 1996) and whether knowledge requires justification (Sartwell 1992); however the issue of whether knowledge requires truth has remained uncontroversial. As Laurence Bonjour remarks, the truth condition is “something almost no philosopher has seriously disputed” (2002: 32). Similarly, John Tienson claims, “it has almost universally been assumed that the requirement of truth is part of the analysis of knowledge” (1974: 289). Even philosophers who think that knowledge cannot be analysed in terms of necessary and sufficient conditions agree that knowledge entails truth (e.g. Williamson 2000: 42).

I characterize this idea as follows:

(Factivity) S knows that p only if p .

Factivity is often justified by appealing to the meaning of ‘knows’ and the uses of sentences of the form ‘ S knows that p ’; for example, non-factive knowledge claims sound odd or contradictory. There seems to be nothing wrong with the sentence, ‘I believe so, but perhaps I am wrong’, however there does seem to be something wrong with the sentence, ‘I know, but perhaps I am wrong’ (Kaplan 2006). Similarly, it seems odd to claim ‘ S knows that p , but not- p ’; indeed, the notion of false knowledge sounds contradictory. In addition, certain combinations of sentences don’t seem to make sense, such as ‘I know that Jess is at home and you know that she is at the office.’² Finally, consider how people withdraw the ascription of knowing that p on the grounds that not- p turns out to be so. In such cases people do (and presumably should) retract the initial claim to know that p by saying something like, ‘Well, I guess I didn’t really know.’ What speakers do *not* tend to do is hold their ground, claiming that one can know that p even though one also believes that not- p .

Allan Hazlett (2010) has recently argued that the concept of knowledge that serves as the meaning of ‘knows’ in ordinary talk is *not* factive.³ In support of this proposal, Hazlett cites a few cases in which ‘knows’ is appropriately used non-factively, where ‘appropriately’ means used in a way that does not strike most people as deviant, improper, unacceptable, necessarily false, etc. Hazlett says that such utterances seem appropriate because they are literally true. In the ordinary sense of ‘knows’ it *is* possible to know false claims; thus, the ordinary concept of knowledge is not factive. If Hazlett is right, a significant upshot is that epistemologists who think that knowledge includes a truth condition should no longer theorize about knowledge on the basis of how ‘knows’ is used. He therefore suggests “a divorce for the linguistic theory of knowledge and traditional epistemology” (*ibid.*: 500).

This paper defends the orthodox view that the ordinary concept of knowledge is factive; thus, an utterance of the form ‘ S knows that p ’ is true only if p . If my argument is successful, it will also undermine Hazlett’s motivation for claiming that epistemologists should not consider linguistic phenomena when theorizing about knowledge. Let us first look at the details of Hazlett’s view.

II. HAZLETT'S NON-FACTIVE THEORY OF 'KNOWS'

In his important article "The Myth of Factive Verbs," Hazlett (2010) cites the following non-factive uses of 'knows', which he claims do not strike ordinary people as deviant, improper, unacceptable, necessarily false, etc.:

- (a) Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers were actually caused by bacterial infection.⁴
- (b) He figures anything big enough to sink the ship they're going to see in time to turn. But the ship's too big, with too small a rudder ... it can't corner worth shit. Everything he knows is wrong.⁵

If these are genuine cases in which 'knows' is appropriately used non-factively, then Factivity faces some recalcitrant data. We need an explanation for why such utterances seem acceptable, and Hazlett suggests that the best working hypothesis is that these utterances are *true*.

Hazlett develops his non-factive theory of 'knows' in the following way. He maintains that while an utterance of the form 'S knows that *p*' might be true even though *p* is false, the following (at least) are necessary conditions on the truth of knowledge ascriptions:

- (NF1)** An utterance of 'S knows *p*' is true only if S believes *p*.
- (NF2)** An utterance of 'S knows *p*' is true only if S possesses epistemic warrant for (her belief that) *p*.

Hazlett appeals to an unorthodox externalist concept of warrant (*ibid*: 508). On his view, there is a plurality of sufficient conditions for epistemic warrant, including two that play a role in his discussion:

- (Proper Source Warrant)** S's belief that *p* is epistemically warranted if S's belief that *p* was formed in a reliable way.
- (Success Warrant)** S's belief that *p* is warranted if S's belief that *p* is true.

Epistemic warrant is understood in a non-internalistic fashion. Hazlett also considers it a conceptual truth that epistemically warranted beliefs tend to be true. Given this concept of knowledge, "the truth of S's belief that *p* is sufficient, but not necessary, for that belief to count as knowledge" (*ibid*: 509).⁶

Hazlett catalogues several facts about our use of 'knows' that a theory of knowledge attributions should be able to explain (*ibid*: 510-11). We only need to focus on the problematic cases for his non-factive theory, namely:

- I. We use 'knows' to relay testimony.** If A and B are police, investigating a recent bombing, for example:

A: Any information from the FBI about how the bomb was constructed?
B: They know the bomb was homemade.

- II. We use 'knows' to describe the beliefs of a third party, while presupposing the truth of the proposition in question.** If A and B are prosecutors, for example:

A: What's relevant is whether the defendant willingly committed a crime.

B: Well, she knew that what she was doing was a crime.

III. We use 'knows' to guarantee the truth of some proposition. If A and B are inspecting a specimen, for example:

A: Can we be sure that this one is of the genus *Calcinus*?

B: I know that this is a specimen of *Calcinus hazletti*. (*ibid.*: 510)

These cases are problematic for the non-factive theory because B suggests that the known proposition is true. In order to resolve this worry, Hazlett appeals to a Gricean (1989) explanation according to which one who utters 'S knows that *p*' typically *implies* that *p* is true. This differs from the traditional idea that 'S knows that *p*' *entails* *p*.

The essentials of the Gricean approach are as follows. People generally assume that their conversational partners are cooperative, which requires mutually assumed conformity to at least three maxims:

Quality: Do not say anything you believe to be false, or which you don't have reason to believe is true.

Quantity: Make your contribution to the conversation as informative, and only as informative, as required.

Relation: Make your contribution to the conversation relevant. (Grice 1989: 26-9)

Let's focus on the aforementioned FBI example in order to outline Hazlett's Gricean explanation. In an illuminating passage, Hazlett writes:

Since it is mutually assumed that speakers are conforming to Quantity and Relation, B here implies that she believes that the bomb was homemade, and that she wishes her interlocutor to believe this as well—for otherwise she would say, of the FBI, that they *think* that the bomb was homemade, but that they are wrong, or something to that effect. To attribute *knowledge* is to say something that entails that the FBI possesses epistemic warrant for their belief that the bomb was homemade. Recall that A is assuming that B will say (and only say) what is relevant. If B thinks that the bomb was not homemade, despite the FBI's warranted belief that it is, then she should not say anything that entails that their belief is warranted, i.e. anything that would misleadingly suggest to A that their belief is *true*, unless she were to explicitly add that their belief isn't true. (2010: 512-3)

Hazlett goes on to compare the FBI case to familiar Gricean-friendly examples, such as the following:

The local who says that there's a gas station around the corner implies that the gas station is open to the public, but 'There's a gas station around the corner' does not entail that the gas station is open to the public. (*ibid.*: 513)

Likewise, A's claim that the FBI know the bomb was homemade implies that the bomb was homemade, but 'They know the bomb was homemade' does not entail that the bomb was homemade, according to Hazlett.

Hazlett mentions two additional points in favor of the non-factive theory. First, "most people do not find the claim that nothing false can be known to be obvious" (*ibid.*: 503), which he takes to show that "as far as common sense goes, the best the defender of

Factivity can say is that common sense doesn't have an opinion about Factivity. At worst, common sense is downright hostile to it" (*ibid*). Second, it is often held that an utterance of the form 'I know p , but not p ' is contradictory, but Hazlett uses the Gricean account to explain the inappropriateness of this type of utterance without appealing to contradiction. Such an utterance is improper because it is paradoxical in the Moorean sense: it is a statement that is not logically contradictory but absurd to utter because it violates certain conversational norms. In particular, one who utters it implies that p is true; thus, the inappropriateness of such an utterance is explained by implication, not entailment. The non-factive theory is thereby made more plausible because it explains the widespread (albeit mistaken) belief that false knowledge is contradictory.

Here's the twist: although Hazlett defends a non-factive theory of 'knows', he nevertheless maintains that truth might be a necessary condition for knowledge. What makes this argument possible is that Hazlett rejects the following linguistic thesis:

(Factivity-L) An (ordinary) utterance of 'S knows that p ' is true only if p .

but he does not reject the following epistemological thesis:

(Factivity) S knows that p only if p .⁷

Factivity-L is a claim about the truth conditions of certain utterances, whereas Factivity is about the necessary conditions for knowledge. This distinction between the linguistic theory of knowledge attributions and what Hazlett calls "traditional epistemology" allows him to reject Factivity-L without sacrificing the idea that knowledge entails truth. Thus, the concept of knowledge that serves as the meaning of 'knows' in ordinary talk is *not* factive, but truth can remain a necessary condition for knowledge.

Hazlett admits that it is often difficult to separate these two theses because Factivity is typically justified by appeal to Factivity-L. This is because we often look at what counts as an appropriate utterance in order to elucidate our concepts. A major upshot of Hazlett's view is that epistemologists will have to look elsewhere to support the idea that knowledge is factive. More significantly, he claims that epistemologists should "stop looking at the linguistic phenomena altogether" (*ibid*: 499). Hazlett endorses this idea because he thinks that traditional epistemology should not be especially interested in the concept of knowledge that serves as the meaning of 'knows' in ordinary talk.

In summary, Hazlett puts forward two challenges.⁸ The first is a *conceptual* challenge against Factivity-L. Hazlett argues that a true utterance of the form 'S knows that p ' does not entail that p is true; thus, the concept of knowledge that serves as the meaning of 'knows' in ordinary talk is not factive. The second is a *methodological* challenge to epistemologists who theorize about knowledge by appealing to knowledge ascriptions and related linguistic phenomena. Hazlett argues that epistemologists can no longer theorize about knowledge on the basis of how 'knows' is used, which targets the trend in recent epistemology to investigate knowledge on the basis of the semantics of 'knows' (e.g., see Cohen 1988; DeRose 1995; Lewis 1999; Schaffer 2004; and Stanley 2005).

It is clear that Hazlett wants to develop a theory of 'knows' suitable for the semantics of ordinary language, however just how this project is connected to what he calls "traditional epistemology" is unclear (although he says that these two projects are not in competition). Should we accept a plurality of concepts of knowledge?⁹ Are epistemologists to look at the language in which philosophical claims are made, as though they mean something different from 'knows' as uttered in everyday talk? Is a theory of 'knows' that is suitable for the semantics of ordinary language of *any* epistemological significance? Hazlett is silent on all of these issues.

Hazlett's methodological challenge succeeds only if we have a reason to think that the relevant statements about the usage of 'knows' sufficiently underdetermine the relevant statements about the nature of knowledge (or 'knowledge').¹⁰ In particular, Hazlett argues that epistemologists have reason to stop looking at the linguistic data *if* they want to keep working on anything like the standard analysis of knowledge, which includes a truth condition.¹¹ Notice, however, that it is unclear why we should remain convinced that knowledge (or 'knowledge') as studied by epistemologists is factive. If Hazlett's non-factive theory of 'knows' is correct, and if philosophers typically justify Factivity by appealing to knowledge ascriptions and related linguistic phenomena, then abandoning Factivity-L substantially weakens the case for Factivity. This is not to suggest that there are no non-linguistically motivated arguments for Factivity; my point is simply that Hazlett does not provide any indication as to where *else* epistemologists might look to justify Factivity now that we should cease examining appropriate utterances as a way of investigating knowledge.

Nevertheless, I do not wish to challenge the idea that knowledge as studied by epistemologists is factive. In other words, I shall grant for the sake of argument that "epistemologists have every right to insist that knowledge (as they understand it) is factive" (Hazlett 2010: 500). Hazlett himself is neutral on whether or not epistemologists should give up Factivity; his claim is that anyone who insists that knowledge (in some epistemologically significant sense) entails truth should not guide their theorizing by how 'knows' is used. In short, any epistemologist who is committed to the factivity of knowledge faces his methodological challenge.

Having presented Hazlett's arguments, including the data he adduces and how they are supposed to support his view, I shall now argue that once we consider a broader range of data it becomes clear that the traditional view is more promising than Hazlett's non-factive theory of 'knows'.

III. A DEFENSE OF ORTHODOXY

An immediate obstacle to assessing the plausibility of Hazlett's theory is that the view, in its present form, makes very few predictions about the relevant cases involving 'knows' because it does not have enough content. For instance, whether in addition to NF1 and NF2 there are other necessary conditions for true knowledge ascriptions is left open. Hazlett also remains neutral on whether NF1 and NF2 are jointly sufficient. If they are, this account would allow for cases in which we may truly ascribe knowledge to subjects who have mere true belief, as well as cases in which we may truly ascribe knowledge to subjects who have reliably formed false beliefs. It is also left open what other conditions suffice for epistemic warrant. For these reasons it is difficult to determine whether Hazlett's account – when sufficiently filled out – will make many false predictions concerning our use of 'knows'. This is a significant obstacle because comparing his theory with the way 'knows' is used in daily life is an indispensable tool for evaluating whether Hazlett achieves his goal of developing a theory of 'knows' that is suitable for the semantics of everyday speech.

This obstacle, though significant, is not insurmountable. The remainder of this paper will demonstrate that considering a broader range of linguistic data provides additional support for the traditional view. First, I will defend the traditional view by providing an alternative account of the propriety of non-factive uses of 'knows'. In particular, I argue that the idea 'protagonist projection' offers a plausible explanation for why 'knows' is sometimes appropriately used non-factively, contrary to what Hazlett argues. Second, I provide some general considerations that favor the orthodox view over the non-factive view. My conclusion is that the traditional view is more plausible than the non-factive theory.

One reason Hazlett thinks we should prefer the non-factive theory to the traditional view is that “charity enjoins us to avoid positing systematic falsehood in ordinary talk” (*ibid.*: 517). All else being equal, this is certainly true; however there is nothing especially worrying about positing falsehood in ordinary talk if the explanation for doing so is plausible. For example, people often say things like

- (c) “It’s 2:30pm” (when it is actually 2:29pm)
- (d) “There’s no coffee left” (when the tin actually contains a few grains)
- (e) “My, aren’t *you* clever” (said mockingly)
- (f) “It’s raining cats and dogs” (when it’s not actually doing so)

We have plausible explanations for why all of these false utterances are appropriate in certain contexts. Take (c) and (d), which seem to be cases of *loose talk*. The ‘loose’ utterance is *close* to being true, it *approximates* a truth. Now, Hazlett rightly points out that the use of ‘knows’ in examples (a) and (b) is “nothing like that” (*ibid.*: 502), since it is not as if the speaker of (a), for instance, was *close enough* to knowing that stress causes ulcers.¹² It also does not seem appropriate to explain the use of ‘knows’ in (a) and (b) by appealing to other common explanations for false but appropriate utterances, such as sarcasm (as in (e)) or metaphor (as in (f)). Are there any other plausible explanations for why ‘knows’ is appropriately used non-factively?

According to Keith DeRose, in many cases it seems that ‘knows’ (or its cognates) is used with a sort of intonation, called ‘focal stress’: e.g., ‘Mary *knew* that she wasn’t going to survive; fortunately, she was wrong’. DeRose takes the need for this special intonation to indicate that the word is not being used literally (2009: 15). Also, these cases often seem acceptable only when the speaker follows the intonated knowledge ascription with an accurate description of the state of affairs that corrects the initial ascription. It seems incorrect for the speaker to flatly assert, ‘Mary knew she wasn’t going to survive’ and then not inform her listeners that Mary actually survived (*ibid.*).

Special intonation may plausibly occur in some cases, but it seems inadequate to explain examples like (a) and (b). In neither case is it obvious that any focal stress is used when the speaker utters ‘knew’ or ‘knows’, nor are these knowledge ascriptions followed by an accurate description of the state of affairs that corrects the initial ascription.

In order to handle Hazlett’s two cases we can appeal to what Richard Holton (1997: 626) calls *protagonist projection*, which is the idea that the speaker adopts the perspective of someone else and describes how the situation appeared from that perspective, rather than providing a true description of the situation. In some cases speakers can do this without intonation, as when speakers are describing what was once (or is elsewhere) commonly *thought* to be known, which the speaker realizes was (or is) in fact false. Holton writes,

I suggest that these sentences work by projecting us into the point of view of the protagonist; let us call the phenomenon *protagonist projection*. In each case the point of view into which we are projected involves a false belief. We describe the false belief using words that the protagonists might use themselves, words that embody their mistake. So we deliberately use words in ways that do not fit the case. (1997: 626)

Thus, supposing there are clear and persuasive non-factive uses of ‘knows’, protagonist projection offers a plausible explanation for why such uses, though appropriate, are not literally true.

Although this idea seems right about many cases, Hazlett doubts that protagonist projection can account for the full range of cases in which a subject may appropriately use ‘knows’ non-factively. Hazlett’s first argument is that speakers who engage in

protagonist projection will respond negatively to questions about whether the situation being described was *really* that way. For example,

the speaker who says ‘He gave her a ring studded with diamonds, but they turned out to be glass’ will respond negatively to the question of whether the ring was really studded with diamonds. It’s not clear how the ordinary person who utters [(a) or (b)] would respond to such a question... (Hazlett 2010: 516)

Hazlett’s argument seems to suppose that if it can be shown that speakers will not respond negatively to such ‘really-questions’, this is evidence for the conclusion that such examples are not cases of protagonist projection.

It is not clear that the linguistic data actually support Hazlett’s view.¹³ Let’s consider the Titanic case (b) in this context. If Hazlett were correct, then the speaker of this statement would not obviously deny that the situation described was *really* that way. But this prediction seems false. Suppose that person A asserts

A: He figures anything big enough to sink the ship they’re going to see in time to turn. But the ship’s too big, with too small a rudder ... it can’t corner worth shit. Everything he knows is wrong.

Now suppose that B asks A

B: So, did he really know that the ship would have time to manoeuvre around any object large enough to see?

How should A respond? A negative answer seems uncontroversially appropriate.

Other examples strengthen this idea. Imagine that A and B are discussing certain beliefs held by ancient peoples.

A: The ancients knew the earth was flat.

B: So, did they really know the shape of the earth?

It seems appropriate for A to immediately retract his statement in response to B’s question. Imagine how improper it would be for A to reply

A: Yes, they really knew the shape of the earth.

Such a reply is certainly significantly less felicitous than a negative response. Hazlett’s first challenge to protagonist projection is therefore unconvincing.

Hazlett’s second argument is that the ‘glass diamonds’ case is disanalogous to the case of ‘knows’. Diamonds are a natural kind, so we have scientific reasons to believe that diamonds are not made of glass; knowledge, in contrast, is not a natural kind—but even if it is, we have nothing approaching *scientific* grounds on which to believe that nothing false can be known. The boundaries of the concept of knowledge are not drawn by nature in the same way that they are for the concept of diamonds. In Hazlett’s own words, the boundaries of the concept of diamonds are not “ours to draw—nature draws them for us” (*ibid.*: 517).

Hazlett is right that there is a difference in the kind of evidence we have for the claim that no diamonds are made of glass than for the thesis that nothing false can be known: we have broadly scientific reasons for believing that diamonds are not made of glass, but not for thinking that nothing false can be known. In this way, the cases are disanalogous. But why should this difference worry us? We do not need anything like scientific grounds

to determine that nothing false can be known because we are dealing with a conceptual truth. Consider that we can establish other obvious conceptual truths, such as the fact that something's being a circle entails that it is not a square, without anything like scientific grounds. Similarly, we do not need scientific reasons for believing that nothing false can be known.¹⁴ The fact that we have independent scientific reason to believe that no diamonds are made of glass is irrelevant to the question of whether the examples we are considering are true or false.¹⁵

Protagonist projection thus offers a plausible explanation for why 'knows' is appropriately used non-factively, contrary to what Hazlett argues.

Another potential problem faced by Hazlett's account, which is not a problem for the traditional theory, is finding clear and persuasive examples to support the view. Hazlett claims that his two examples do not strike ordinary people as deviant, improper, unacceptable, necessarily false, etc., and he no doubt thinks that such cases can be multiplied. Unfortunately he provides no empirical evidence to support these claims. In the absence of such evidence, one worry is whether many people actually do have the intuitions Hazlett claims they do. For what it's worth (and it's probably not much, since Hazlett might be able to confirm his own intuitions in the same way), the examples he cites *do* strike everyone I have asked as improper. But perhaps this merely illustrates that Hazlett and I keep different company.

I want to avoid an intuition stalemate. The best way to resolve this issue would be through experimental work, namely conducting a study with carefully constructed vignettes assessed by a representative sample of the population; however, this task lies beyond the scope of this paper. But even if acceptable non-factive uses of 'knows' are acknowledged, support for the traditional view can still flow from competent speakers' inclinations to accept all instances of the schema:

(I) If 'I know that p ' is true, then p .

Or negative answers to all instances of questions of the form

(II) Does S know that p even though p is false?

It strikes me as plausible that all competent speakers are so inclined. Consider the following instance of (I), with which everyone should agree: If 'I know that Tom is in France' is true, then Tom is in France. Also, my discussion of 'really-questions' suggests that speakers would be inclined to accept negative answers to the relevant instances of questions with form (II).

Hazlett also tries to motivate his position by appealing to the fact (if it is a fact) that "most people do not find the claim that nothing false can be known to be obvious" (*ibid.*: 503). That 'knows' is not *obviously* factive to some competent speakers is clearly compatible with it being factive. (A parallel: few people may find 'Everything you know you believe' to be obvious, yet it is commonly thought that 'knows' entails belief.) Hazlett is not arguing that 'knows' isn't factive because many people do not find the claim 'nothing false can be known' to be obvious. Rather, his claim is that *common sense* does not support Factivity-L.

He again cites no evidence to support this claim, but let's set that aside. I do not wish to rest my case on the assumption that Hazlett is wrong about what people find intuitive or acceptable. Thus, I shall henceforth assume for the sake of argument that the empirical data will confirm Hazlett's claim that a significant number of people find many non-factive utterances acceptable.

Even if we presume that most people do not find the claim ‘nothing false can be known’ to be obvious, nothing prevents us from rephrasing this point in ways that people *do* find intuitive.¹⁶ For example, ask someone ‘Did the ancients know the earth was flat, even though it isn’t?’ or ‘Do some children know that Santa Claus exists, even though he doesn’t?’ and their responses will indicate that they find the idea of false knowledge odd, if not impossible. The linguistic data in support of this point can easily be multiplied. For this reason Hazlett seems to drastically overstate his case by claiming that “as far as common sense goes... [it] doesn’t have an opinion about Factivity[-L]. At worst, common sense is downright hostile to it” (*ibid.*: 503). There is much linguistic data in support of Factivity-L, which suggests that far from being “hostile” to this thesis, common sense overwhelmingly favors it.

Recall that Hazlett prefers the non-factive theory to the traditional view because he thinks it better avoids positing systematic falsehood in ordinary talk. Yet positing falsehood in cases where ‘knows’ is appropriately used non-factively would only be a strike against the traditional view if the non-factive theory did comparatively better on this score—which it doesn’t. The traditional view must posit falsehoods in those relatively few cases in which ‘knows’ is deliberately used non-factively, whereas the non-factive view seems to posit a large number of falsehoods in everyday speech because it commits speakers to the widespread mistaken usage of ‘knows’. There are many problematic cases involving knowledge denials, for example, as when people claim they ‘thought they knew’, which implies that they *didn’t* know, despite being warranted in their belief because it was formed in a reliable way. Suppose, for instance, that A and B intend to go for a walk and that A proposes that they walk in Cascadilla Gorge.¹⁷ B protests that he would like to walk beside a flowing stream and that at this time of year the gorge is probably dry. In response, A says

A: I know that it won’t be dry because I saw a lot of water flowing in the gorge when I passed it this morning.

Presume that A walks by the gorge every day and has done so for several years. Also presume that on every past occasion in which he has observed a lot of water flowing in the gorge in the morning, he has also observed water in the evening. If A and B go and find water, we would not hesitate to say that A knew. Imagine, however, that they find the gorge to be dry. We should not say that A knew, but rather that he *thought he knew*. Indeed, A himself likely would – and presumable should – say some such thing.

Likewise, there are cases in which knowledge is denied on the grounds that the subject ‘only thinks that *p*’. Imagine that Barry Marshall and Robin Warren have just proven that ulcers are actually caused by bacterial infection. Before they share this news with their colleagues, they overhear two physicians discussing the adverse effects of stress. One of them says,

Physician: I know that stress causes many problems, including ulcers.

Having heard this, Barry says to Robin,

Barry: Ha! He *only thinks* that stress causes ulcers. Let’s go tell him about our results.

On Hazlett’s view, we may be systematically mistaken for denying knowledge in such cases because it is permissible to ‘know’ something false. Admittedly, it is difficult to say for sure whether these judgments are mistaken on Hazlett’s view because it currently lacks enough content. Nevertheless, at the very least we can say that Hazlett has not

given us any reason to think that we do not fall prey systematically to error in these cases. After all, all the necessary conditions for the truth of the relevant knowledge attributions he offers (belief and warrant) are fulfilled in these cases and his account doesn't contain anything else that would allow him to secure the desired results here—this is bad news for Hazlett. What's even worse is that it is hard to see how he could expand his account (in a non-*ad hoc* way) so as to get the desired results. It does not come as a surprise, then, that Hazlett himself acknowledges that these kinds of cases are problematic for his account (*ibid.* 510).¹⁸

In contrast, the traditional view neatly explains all of the linguistic data, such as why people deny or retract knowledge ascriptions when the thing 'known' turns out to be false (because 'knows' is factive), why 'knows' is sometimes appropriately used non-factively (because of protagonist projection), why 'knows' is used to guarantee the truth of some proposition (because it entails truth), and why so many non-factive uses of 'knows' sound contradictory (because they are). Thus, although Hazlett is right that we should prefer a theory that avoids positing systematic falsehood in ordinary talk, his own view does not enjoy this advantage.

John Turri (2011) identifies two more problems for the non-factive theory of knowledge, which I shall briefly mention. First, Hazlett's view presumes that conversational implicatures are directly and felicitously cancellable, yet 'knows' seems to fail this test. To illustrate this point, consider B's claim that the FBI "know the bomb was homemade". According to Hazlett, B implies that the FBI's belief is true by ascribing knowledge to the FBI without explicitly cancelling the suggestion that what the FBI knows is true. This is supposed to be compatible with the familiar Gricean idea that implications can be directly and felicitously cancelled. For example, the local who says that there is a gas station around the corner implies that the gas station is open to the public, unless he cancels the implication—i.e., 'There's a gas station around the corner, but it's been closed for months'. In this case directly cancelling the implicature sounds perfectly acceptable and therefore presents no problem. Notice, however, that it would sound improper if B were to directly cancel the implication from 'They know the bomb was homemade' to 'the bomb was homemade'. Suppose B had said,

B: They know the bomb was homemade, but the bomb wasn't homemade.

This statement is highly counterintuitive. Hazlett's account therefore conflicts with the familiar Gricean idea that implications can be directly and felicitously cancelled. Grice (1989: 44) thinks that all conversational implicatures are cancellable, and a cancellability test is frequently used to show that some effect of language use is not an implicature (Borge 2009).¹⁹

Second, Hazlett's preference for the literal reading of (a) is problematic given that it is already clearly a case of exaggeration: 'Everyone knew that stress caused ulcers' doesn't literally mean *everyone*, nor does it mean that stress was known to cause *all* ulcers (Turri 2011: 146). It therefore seems acceptable to dismiss (a) as an overstatement, beginning with the obviously false claim about everyone, and continuing with the attribution of knowledge. Similarly, it is obvious that not *everything* the speaker in (b) knows is false.

Thus far I have argued that Hazlett's non-factive theory of 'knows' faces several unacknowledged problems and that his objections to the traditional view are unconvincing. However I have not gone so far as to suggest that the traditional view can handle *all* of the linguistic data. It might be true that there are some (very few) clear and persuasive cases in which 'knows' is used non-factively, with apparent propriety, and that such cases cannot be explained away by appealing to special intonation, protagonist projection, or some other hypothesis. Even presuming that there are such cases (it has

not been shown that there are), I shall now argue that it is still not obvious that we should find them compelling.

The proposal I shall sketch is as follows: presuming we are left with some recalcitrant data, we can nevertheless reject such data in order to have a concept that is suitably 'cleaned up' in the sense that it is the product of theory. This requires some account that identifies the central cases as central, but I see no *prima facie* reason to doubt the prospects of such an account. This proposal is motivated by the idea that a theory of 'knows' should in large part be judged by its ability to explain core features of our knowledge-ascribing practices, but it need not slavishly yield to the vagaries of ordinary usage. There might be a margin of error between the analysed concept and everyday speech, so we should not insist that our account must coincide perfectly with our linguistic intuitions. This parallels the widely held view that our considered judgments are subject to occasional irregularities, inconsistencies, and distortions; thus, when presented with a theory that gives an appealing account of those judgments, we may wish to revise some of them to conform with it (Shaw 1980: 129).

This proposal may be understood in several ways.²⁰ One way of characterizing this idea is by the model of reflective equilibrium. This methodological practice tests theories against our intuitions about cases, with the goal of finding consistency between theory and intuition. However not all intuitions are authoritative. According to John Rawls, to whom we owe the name 'reflective equilibrium', only intuitions that have been appropriately filtered are part of the set of intuitions with which equilibrium must be struck. In particular, a plausible theory must be consistent with those intuitions that are both widely shared and firmly held after careful consideration (Rawls 1971: 47-8). I worry that none of the data Hazlett cites satisfy these criteria and thus they are not part of the base of intuitions with which equilibrium must be struck;²¹ however I shall set aside this reservation and assume that his examples are widely held upon careful reflection. Even so, many advocates of reflective equilibrium grant that on some occasions we may even decide to give up some of our *firmest* intuitions for the sake of theoretical unity. Nelson Goodman, who first proposed the reflective equilibrium test, seemed to hold this view (1965: 66). Thus, although a theory passes the reflective equilibrium test only if it agrees with our widely shared and firmly held intuitions, it need not agree with *all* such intuitions.

A second way of characterizing this proposal is by analogy with the scientific method. There is a long tradition of scientists ignoring recalcitrant data if its percentage is insignificant and if a theory seems otherwise flawless (Rollin 2006: 258). In 'The Role of Anomalous Data in Knowledge Acquisition', for instance, Chinn and Brewer postulate that "there are seven distinct forms of response to anomalous data, only one of which is to accept the data and change theories. The other six responses involve discounting the data in various ways in order to protect the preinstructional theory" (Chinn and Brewer 1993: 1). Data discounting occurs especially in cases where human error is likely involved; so if we reasonably presume that speakers' linguistic intuitions are subject to occasional irregularities, inconsistencies, and distortion, then we may wish to ignore such intuitions if they are anomalous. We may defend the traditional view by responding to recalcitrant data in the same way that scientists do.

A third way of understanding this proposal is as an explication of 'knows'. There is disagreement as to the precise nature and purpose of the brand of philosophical analysis variously called 'explication' (see Hanna 1968); however it is generally agreed that this method consists in replacing a vague, pre-systematic notion (the *explicandum*) with a more precise notion (the *explicatum*). The central feature of explication is a *correspondence condition*, which places a considerable degree of emphasis on the similarity between the pre-systematic concept and the reconstructed concept. This condition permits the

reconstructed concept to deviate from our ordinary language conventions but also requires it to respect the application of the pre-systematic term. The aim of explication is not to replace a familiar notion with a novel one; on the contrary, the aim is to capture the actual meaning of our familiar notion. Explication is motivated by the idea that some of our language conventions are possibly inconsistent and that we may remedy this defect by tidying up the ordinary concept. We begin by analysing our familiar notion of 'knows' and then propose modifications in an effort to tidy up the notion. So long as the proposed changes are motivated by considerations of simplicity and do not result in any radical departures from the familiar concept that serves as the meaning of 'knows' in ordinary talk, it counts as an acceptable form of analysis (Stich 1988: 33).

All three of these approaches have in common the idea that our analysis should be anchored in everyday usage but not hostage to the linguistic data in an overly austere way. These approaches provide reason to favor the traditional view over the non-factive view due to considerations having to do with simplicity and range of acceptable attributions/denials explained.

One might worry that there is something suspiciously *ad hoc* about this move. Am I not simply ignoring data that conflict with my view? This is not my suggestion. My claim is that once we take all the data into account we adopt the theory that does the best job of handling most of the data, presuming that no account gets all of it right. After all, we need not (and probably should not) treat *everything* people say as unalterable data; the way people speak can sometimes be misleading.²² Thus, the awkward linguistic data can perhaps be ignored as aberrations. Even J.L. Austin, the figurehead of ordinary language philosophy, held that facts about language are the first word in philosophy, but not the last word about how we are to think about a concept. Ordinary language can be supplemented and improved upon and superseded (1970 :185).

The situation I have described is not ideal. A theory of 'knows' that encounters no recalcitrant data would be preferable; unfortunately, the possibility that such a theory exists is very slim. The next best theoretical option is to supplement our theory of 'knows' with a plausible explanation that handles recalcitrant data, such as protagonist projection or special intonation. I have argued that these hypotheses provide plausible explanations for the examples we have encountered. In order to be charitable, however, I have also presumed that there may be cases in which 'knows' is used non-factively, with apparent propriety, and that such cases cannot be explained away by appealing to special intonation, protagonist projection, or any other hypothesis. Nevertheless, even if we grant this assumption I have provided some general theoretical reasons for legitimately discounting these problematic cases.²³

IV. CONCLUSION

"The Myth of Factive Verbs" successfully shows that the connection between knowledge (or 'knowledge') and truth is not as clear-cut as it is often assumed. Does Hazlett also succeed in turning the tide in favor of the non-factive theory? This paper has argued that the case for the orthodox view remains strong.

I'll conclude by pointing out that even if Hazlett is correct, a consequence of his view is that we can still *know* that the ordinary concept of knowledge is factive. Factivity-L may be known because this is compatible with Factivity-L being false. Faced with this fact, Hazlett might claim that *I* know that 'knows' is factive, but *he* knows that it isn't; or perhaps he would claim to know that 'knows' is factive and also that this very assertion is false. If he adopts either of these positions, he is committed to the view that mutually inconsistent things can be 'known'.²⁴ Hazlett might be happy with this consequence; I consider it yet another strike against his view.²⁵

References

- Achenbach, Joel. 2005. "Cat Carrier: Your cat could make you crazy." *National Geographic* 208.
- Austin, John L. 1970. *Philosophical Papers*. Oxford: Oxford University Press.
- Bonjour, Laurence. 2002. *Epistemology : Classical Problems and Contemporary Responses*. Lanham MD: Rowman and Littlefield.
- Borge, Steffen. 2009. "Conversational Implicatures and Cancellability." *Acta Analytica* 24 (2): 149-154.
- Chinn, Clark & Brewer, William. 1993. "The Role of Anomalous Data in Knowledge Acquisition: A theoretical framework and implications for science instruction." *Review of Educational Research* 63: 1-49.
- Cohen, Stewart. 1988. "How to Be a Fallibilist." *Philosophical Perspectives* 2: 91-123.
- Davis, Wayne. 2007. "Knowledge Claims and Context: Loose Use." *Philosophical Studies* 132 (3): 395-438.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104 (1): 1-52.
- , 2009. *The Case for Contextualism*. Oxford: Oxford University Press.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121-123.
- Grice, Herbert P. 1989. *Studies in the Way of Words*. Boston: Harvard University Press.
- Hazlett, Allan. 2010. "The Myth of Factive Verbs." *Philosophy and Phenomenological Research* 80 (3): 497-522.
- Holton, Richard. 1997. "Some Telling Examples: A Reply to Tsohatzidis." *Journal of Pragmatics* 28: 625-628.
- Kaplan, Mark. 2006. "If You Know You Can't Be Wrong." In Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, 180-98.
- Kornblith, Hilary. 2002. *Knowledge and its Place in Nature*. Oxford: Oxford University Press.
- Lewis, David. 1999. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549-567.
- Malcolm, Norman. 1952. "Knowledge and Belief." *Mind* 61 (242): 178-189.
- Radford, Colin. 1966. "Knowledge: By Examples." *Analysis* 27 (1): 1-11.
- Rollin, Bernard E. 2006. *Science and Ethics*. Cambridge: Cambridge University Press.
- Sartwell, Crispin. 1992. "Why Knowledge is Merely True Belief." *The Journal of Philosophy* 89 (4): 167-180.
- Schaffer, Jonathan. 2004. "From Contextualism to Contrastivism." *Philosophical Studies* 119 (1-2): 73-104.
- Shaw, William. 1980. "Intuition and Moral Philosophy." *American Philosophical Quarterly* 17: 127-34.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stich, Stephen. 1988. "Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity." *Synthese* 74 (3): 391-413.
- Stokke, Andreas. *Forthcoming*. "Protagonist Projection." *Mind and Language*.
- Tienson, John. 1974. "On Analysing Knowledge." *Philosophical Studies* 25 (4): 289-293.
- Turri, John. 2011. "Mythology of the Factive." *Logos and Episteme* 2 (1): 143-152.
- Weiner, Matt. 2009. "The (Mostly Harmless) Inconsistency of Knowledge Attributions." *Philosophers' Imprint* 9 (1): 1-10.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

¹ Although the truth condition is almost universally accepted, philosophers have offered some opposing ideas about what it is for a belief to be true.

² Compare 'I believe that Jess is at home and you believe that she is at the office', which sounds perfectly fine.

³ Hazlett suggests that several two-place predicates relating a person to a proposition are not factive, such as ‘learns’, ‘remembers’, and ‘realizes’; however he focuses specifically on ‘knows’, as I shall.

⁴ Adapted from Achenbach 2005.

⁵ Brock to Bodine, in *Titanic* 1997.

⁶ This claim by Hazlett is actually misleading. Hazlett does not commit to the view that true belief is sufficient for a statement of the form ‘S knows that *p*’ to be true—although this *is* perfectly consistent with his view. Rather, the idea is that the truth of S’s belief that *p* is sufficient, but not necessary, to confer epistemic warrant for (her belief that) *p*, and that such epistemic warrant is a necessary condition on the truth of knowledge ascriptions. Thanks to an anonymous reviewer for pressing me to clarify this point.

⁷ In his own article, Hazlett refers to what I call ‘Factivity’ as “the truth condition” and what I call ‘Factivity-L’ as “Factivity”. See Hazlett 2010: 499.

⁸ Turri 2011 identifies the same two challenges.

⁹ Hazlett shows some sympathy for this view near the end of his paper. Notice, however, that if we accept a plurality of concepts of knowledge, such that there is (at least) an ordinary non-factive concept and a factive concept studied by epistemologists, then this might be the *only* difference between these concepts. If this were so, then epistemologists would do well to study the ordinary concept carefully to elucidate the epistemologically significant concept of knowledge. Turri 2011: 145 makes a similar point. This would undermine Hazlett’s methodological challenge.

¹⁰ Hazlett is neutral on the question of whether epistemologists should investigate the nature of knowledge or of some *concept* of knowledge. See Hazlett 2010: 499.

¹¹ An anonymous reviewer rightly pointed out that Hazlett’s methodological challenge does not require the assumption that knowledge is factive. Rather, *any* reason to think that our usage of ‘knows’ sufficiently underdetermines claims about the nature of knowledge will be adequate to motivate his methodological challenge. While this point is surely correct, the only reason Hazlett provides us to think that ‘knows’ sufficiently underdetermines statements about the nature of knowledge is that the latter is presumable factive.

¹² Wayne Davis 2007 defends the view that ‘S knows *p*’ is commonly used loosely to implicate ‘S is close enough to knowing *p* for contextually indicated purposes.’

¹³ A similar argument is made by Andreas Stokke in a forthcoming paper. Thanks to Jennifer Nagel for bringing this to my attention.

¹⁴ There is much more to be said here about the nature of concepts and how we discover them, but this goes beyond the scope of my paper, which aims to respond to Hazlett’s arguments. It would be nice if a particular view about the nature of concepts was on the table, but Hazlett offers none.

¹⁵ Hazlett has another objection to the protagonist projection view; namely that cases of protagonist projection require the false utterances to be *deliberate*—the speaker must know that her utterance is false. He claims it is not as clearly plausible to suppose that speakers who appropriately use ‘knows’ non-factively believe that their utterance is false. I have already provided the resources to respond to this objection. Recall the recently argued claim that speakers who use ‘knows’ non-factively tend to deny that the situation described was “really” that way. If this is correct, then it seems to demonstrate that ‘knows’ is used in a way that speakers recognize is false, since they will deny that the situation was really (i.e. truly) that way.

¹⁶ My suspicion is that the reason people do not find this statement obvious is due to phrasing.

¹⁷ I borrow this case from Norman Malcolm’s 1952: 178-189.

¹⁸ Here, I am indebted to the comments of an anonymous referee.

¹⁹ According to Timothy Williamson 2000: 35, “If Φ is an FMSO [factive mental stative operator, such as ‘knows’], the implication from ‘S Φ s that A’ to ‘A’ is not cancelable.” Also, Hazlett himself says “Cancelability is a mark of conversational implicature” in his earlier paper ‘Knowledge and Conversation’ 2009: 597. See Weiner 2006 for a contrasting view and Blome-Tillman 2008 for a reply.

²⁰ Thanks to Rob Trueman for a helpful discussion on these points.

²¹ John Turri puts this point well: “If anything serious hinged on the truth of this [non-factive] statement, I would object: “You don’t *really* mean they knew. You mean that they *thought* they knew, or some such thing.” If the person insisted that he “literally meant” exactly what he said, I would question whether he was using ‘literally’ literally. Supposing he answered affirmatively and I believed him, I would conclude that he was incompetent or confused.” See 2011: 147.

²² The obviousness of this view is easily demonstrated by considering its alternative: a view according to which the way we speak must be taken as absolute, conclusive data. Such a view is clearly misguided, for it entails that speakers can never commit errors in speech. The appropriate stance toward our linguistic intuitions must involve *some* sort of error theory, otherwise we must think we can *never* be mistaken.

²³ One might instead argue that ‘knows’ is polysemous, i.e., in some contexts it takes a factive sense and in others a non-factive sense. Following Hazlett (who follows Grice), I assume that positing polysemy is a

vice, *ceteris parabus*, for a linguistic theory. I have therefore assumed that ‘knows’ is univocal and that a consistent semantics can be given for it. Weiner 2009 disagrees.

²⁴ Hazlett might accept some consistency principle for ‘knowledge’; e.g., for all subjects S, propositions *p*, and times *t*, if S ‘knows’ that *p* at *t*, then there is no subject R such that R ‘knows’ that not-*p* at *t*. However such a principle seems incompatible with Hazlett’s titanic case.

²⁵ Thanks to Hallvard Lillehammer, Jess Kwong, and a very helpful anonymous referee of this journal for valuable advice. This paper was partially funded by the Social Sciences and Humanities Research Council of Canada.