Factive Verbs and Protagonist Projection

Document Version
Accepted author manuscript

Citation for published version (APA):
Buckwalter, W. (2014). Factive Verbs and Protagonist Projection. *Episteme* , *11*(4), 391-409.
https://www.cambridge.org/core/journals/episteme/article/factive-verbs-and-protagonist-projection/C8AAA7F55068A37ECFC7A41E6CE88602

Published in:
Episteme

Citing this paper
Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

Factive Verbs and Protagonist Projection[*]

Wesley Buckwalter

Nearly all philosophers agree that only true things can be known. But does this principle reflect actual patterns of ordinary usage? Several examples in ordinary language seem to show that 'know' is literally used non-factively. By contrast, this paper reports five experiments utilizing explicit paraphrasing tasks, which suggest that these non-factive uses are actually not literal. Instead, they are better explained by a phenomenon known as protagonist projection. It is argued that armchair philosophical orthodoxy regarding the truth requirement for knowledge withstands current empirical scrutiny.

Keywords: knowledge, truth, factivity, ordinary language, projection

1   Introduction

Philosophical debate over the nature of knowledge is an uneven affair. For long stretches, some purported requirements of knowledge were unquestionably assumed by all. For instance, over the past several decades it has been assumed that knowledge must be reliably produced (Dretske 1981; Feldman 2003; Goldman 1986; Nozick 1981; Russell 1948 [2012]; Sellars 1963; Sosa 1991). For centuries it has been assumed that knowledge requires belief (Armstrong 1969; Chisholm 1989; Cohen 1966; Jones 1971; Lehrer 1974). And for eons it has been assumed that knowledge requires truth (Plato [1997] 186c-187b).

---

[*] This is the penultimate version of a paper to appear in *Episteme*. Please cite the final, published version if possible.

Philosophers have also long strived to develop theories that respect ordinary practice (Aristotle [1984]; Moore 1942; Reid 1785 [2002]). In modern times, this impulse is most clearly represented in ordinary language philosophy (Austin 1956). Following Austin, it is widely accepted that patterns of ordinary usage place important constraints on philosophical theorizing about knowledge (Cohen 1999; DeRose 2009; Hawthorne 2004; Fantl and McGrath 2009; Rysiew 2001; Stanley 2005; Stroud 1984; Vogel 1990). On this view ordinary ascription practices serve as powerful evidence in favor of a theory of knowledge. Running roughshod over ordinary practice constitutes grounds for rejection or revision.

If we want a theory of knowledge that simultaneously illuminates its subject matter and respects ordinary practice, we need to know what ordinary practice is like. In recent years, this has inspired collaborative research at the intersection of philosophy and cognitive science (Friedman & Turri in press; Nagel, San Juan & Mar 2013; Starmans & Friedman 2012; Turri, Buckwalter & Blouw in press). The goal of this research has been to empirically investigate knowledge attributions and the psychological mechanisms that generate them (for brief overviews see Beebe 2014; Buckwalter 2012).

In some cases the empirical research has vindicated orthodox assumptions about knowledge. It has completely overturned them in others. For instance, empirical work has challenged the assumption that knowledge must be reliably produced (Sackris & Beebe 2014; Turri, Buckwalter & Blouw in press). On the other hand, researchers have offered powerful folk psychological evidence that knowledge requires belief (Buckwalter, Rose & Turri in press; Rose & Schaffer 2013). Patterns of ordinary usage indicate a firm

commitment to the entailment thesis, even if early results appeared to have questioned it (Murray et al. 2013; Myers-Shulz & Schwitzgebel 2013).

This research has led philosophers and cognitive scientists to debate whether knowledge requires justification and, lately, whether it requires belief. But there has been virtually no controversy over whether knowledge requires truth. Nearly all epistemologists agree that a person can know a proposition only if that proposition is true. As one commentator puts it, "it is overwhelmingly clear that what is false cannot be known" (Steup 2006).

However these days are now over too. In a pair of papers, "The Myth of Factive Verbs," and "Factive Presupposition and the Truth Condition on Knowledge," Allan Hazlett presents compelling evidence against the truth requirement for knowledge (2010, 2012). He gives several ordinary examples of 'knows' used in sentences of the form 'S knows that p' where p is false. If these examples are interpreted literally, they suggest that ordinary speakers ascribe knowledge in the absence of truth. Far from finding this thesis "overwhelmingly clear," such patterns of ordinary usage appear to reveal a blatant disregard for this requirement.

The goal of this paper is to evaluate this very surprising development. Before challenging orthodoxy we need to examine the linguistic evidence more carefully. There have been many studies on knowledge ascription, but none to date have directly focused on the truth requirement, in light of Hazlett's arguments and data. Will controlled experimentation reveal that ordinary language undermines the truth requirement or vindicates it? To anticipate the conclusion, I find that the truth requirement, the most orthodox of orthodox commitments in epistemology, survives largely intact. Using prior

proposals in the literature as a starting point, five experiments are presented to illuminate the relationship between knowledge and truth in ordinary judgments. Results suggest that the putative non-factive uses of 'knows' to be considered are not interpreted as literal ascriptions of knowledge. Alternatively, results suggest that they are better explained by a phenomenon known as 'protagonist projection' (see Holton 1997).

In what follows, Section 2 reviews the received linguistic evidence against factivity presented in the literature, and introduces protagonist projection as a rival explanatory hypothesis to these data. Section 3 argues that the linguistic evidence presented against factivity to date is better explained by the tendency to adopt the perspectives of putative knowers, and not because people think that subjects with false beliefs have knowledge. Sections 4-5 discuss these results and some methodological implications for ordinary language approaches in epistemology.

2   Against Factivity

The orthodox assumption just about unanimously accepted by philosophers is that only true things can be known. Call this the standard truth condition of knowledge:

STC: S knows that p only if p is true.

If patterns in ordinary language support the standard truth condition, then the word 'know', much like the English words 'learn', 'remember' and 'realize' is factive.[1] Call this the Factivity principle:

> *Factivity*: Certain two-place predicates, including 'knows', 'learns', 'remembers', and 'realizes', which denote relations between persons and propositions, are factive in this sense: an utterance of 'S knows p' is true only if p, an utterance of 'S learned p' is true only if p, and so on. (Hazlett 2010, 499)

Our task will be to investigate whether ordinary English uses of these words—and in particular the word 'know' in utterances of the form 'S knows that p'—truly reflect this principle.

From the start, the prognosis for the principle seems grim. A quick Google search reveals that one is surrounded by what appear to be blatantly non-factive uses in ordinary language (headline: "The Cuban missile crisis: what you know is wrong").[2] Consider four more examples—taken from actual press releases, movies, and newspapers—ordinary speakers presumably find acceptable instances of non-factive use. These items are originally presented in Hazlett (2010), and remain the best linguistic evidence to date for non-factivity:

---

[1] For more on factivity in the linguistics literature see Kiparsky & Kiparsky (1971). The notion of Factivity presently adopted is restrictive in the sense that it assumes that predicates like 'knows' involve relations to propositions.

[2] Taken from The New York Times website:

<http://www.nytimes.com/2008/06/13/opinion/13iht-edbeam.1.13693498.html>

Know Ulcers: Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.

Know Sailing: 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.

Learn History: 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.

Realize Death: I had trouble breathing, sharp pains in my side, several broken ribs and a partially collapsed lung, and I was in the middle of nowhere without any real rescue assets – it was then that I realized I was going to die out there. (Hazlett 2010, 501, labels added)

The observation that such sentences appear acceptable is an important first start toward helping approximate ordinary usage and knowledge ascription. At the same time, we do not know whether or not the acceptability of these uses stems from the fact that people think that 'S knows p' but not-p is true, or if their acceptability can be accounted for in some other way that does not call into question this requirement for knowledge.[3]

---

[3] A further empirical question yet to be explored is the frequency of these sorts of non-factive uses compared to clearly factive uses. This paper remains neutral on this point, focusing instead on providing an explanation for such constructions when they arise (though see Turri 2011b and Tsohatzidis 2012).

Reflecting on this question has inspired some philosophers to begin exploring putative cases of non-factive usage and how traditional epistemic theorizing might possibly accommodate them. John Turri, for instance, advances a performance-view of knowledge that allows for knowledge of "approximate truths" (2011a; forthcoming). Potential moves non-factive theorists might make have been explored in the work of Daniel Nolan (2008). The possibility that epistemic contextualism might allow for contexts in which certain kinds of false beliefs qualify as knowledge has been noted in the works of Keith DeRose (2009).

By far the most comprehensive response to date has been given by Hazlett. Hazlett (2010) rejects Factivity outright, and argues that the ordinary concept serving as the meaning for 'knows' in ordinary language allows for knowledge under false belief. In place of orthodoxy, Hazlett eliminates the truth condition and argues that the ordinary concept of knowledge is based on the following two necessary conditions:

NF1: An utterance of 'S knows that p' is true only if S believes p.

NF2: An utterance of 'S knows that p' is true only if S possess epistemic warrant for (her belief that) p.

It is an open empirical question whether this kind of non-factive account better characterizes the ordinary concept over a more traditional account of knowledge. Providing for the variety of different ways that 'know' is used in ordinary talk is one indication.

One upshot of an account based on something like NF1 and NF2 is that it is better able to explain the prima facie counterexamples to Factivity considered above. For example, on this view the reason why 'everyone knew that stress causes ulcers' is an

acceptable sentence is because—given a belief backed by sufficient epistemic warrant—people really did have knowledge that stress causes peptic ulcers before it was proven to be H. pylori (Marshall et al. 1985).[4]

One downside of such an account is that the standard truth condition is justified at the cost of ordinary language. The analysis of knowledge that philosophers have historically given is factive through and through. We might not expect the traditional philosophical characterization to perfectly capture every and all ordinary conceptions of knowledge. But it would be very shocking to discover that the traditional philosophical conception did not track at least one major conception, predominately used by ordinary people under the same term. Moreover, if predominate use of the ordinary term is completely at odds with a factive account, it is unclear what use epistemologists would have for ordinary language in the first place (or perhaps even more worrisome, what use ordinary speakers would have for professional epistemologists). Thus it would be extremely troubling to discover that the concept of knowledge informing the definition of 'know' in ordinary language departs so radically from the things philosophers have said about knowledge.

---

[4] For a fascinating account of this discovery (in which Marshall describes infecting himself to prove his theories correct against overwhelming medical consensus) see the Nobel Prize speech available at <http://www.nobelprize.org/mediaplayer/index.php?id=614>. Ironically, Marshall begins with a quote from historian Daniel Boorstein, "The greatest obstacle to knowledge is not ignorance, it is the illusion of knowledge."

It is worth pausing here to briefly consider two repercussions of this eventuality. One the one hand, if we accept ordinary language as an important constraint on epistemic theorizing, then the orthodox characterization of the knowledge relation could very well belong in the waste can. On the other hand, if we retain the standard truth condition and jettison ordinary language constraints, then a significant portion of what motivated and provided evidence supporting leading epistemic theories also relying on ordinary language evidence (e.g. DeRose, Stanley, Fantl & McGrath, and others) may too be lost to us (we return to this point in Section 4).[5]

Given these theoretical costs a more thorough investigation of the linguistic data above is warranted. After all, there could be many other interpretations of the linguistic data that do not challenge Factivity or raise questions about the evidentiary value of ordinary ascriptions. Another hypothesis worthy of consideration is that the sentences presented above do not constitute literal ascriptions of knowledge. One way in particular in which such sentences might not constitute literal ascriptions involves the phenomenon of protagonist projection proposed by Richard Holton:

> 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

---

[5] This is an initial presentation for the purposes of motivating the discussion. Of course, some philosophers may be more willing to jettison ordinary language constraints over others but a full discussion of the merits and costs of this approach must be saved for another occasion.

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. (Holton, 1997, 626)

Some sentences invoke projective readings. For instance, Holton considers the following examples, 'She sold him a pig in a bag. When he got home he discovered it was really a cat,' and 'I saw a shooting star last night. I wished on it, but it was just a satellite.' Sentences like these are acceptable to us because presumably, we take up the protagonist's perspective and then imagine what seems true from their point of view. Since the protagonist's perspective involves a false belief in each case, we use words the protagonists might use to describe those beliefs from their own points of view. So when we read sentences like the examples above, we don't take them literally as evidence that some pigs are actually cats, or that some satellites are actually meteoroids. Rather, we find the sentences acceptable because from the protagonist's point of view, 'he thought she sold him a pig,' or 'she thought she saw a shooting star.' Another way to think about this phenomenon is that it is roughly equivalent to when one talks about someone else, and—with sufficient cues, e.g. imitating their bodily language and tone of voice—can then use 'I' to refer to this other person.[6]

Might this also be happening in the sentences taken to challenge Factivity above? Consider for instance, Know Ulcers. It could be that 'everyone knew that stress caused

---

[6] Similarly, see DeRose (2009) for an argument that special intonations of 'knows' provide crucial cues that sentences implying that protagonists know things that are false are not meant to be taken literally.

ulcers' does not literally count as evidence that 'know' is non-factive, in just the same way that 'I saw a shooting star last night' does not literally provide evidence that some satellites are meteoroids. Instead, perhaps people think Know Ulcers is acceptable only because they assume the protagonist's perspective, in this case the perspective of putative knowers before Marshall. Then, they describe how the situation may have appeared from that perspective. People living before the 1980s had relatively decent evidence supporting their beliefs about the relationship between stress and ulcers, in a way typically thought to be consistent with knowledge. Therefore 'everyone knew that stress caused ulcers' is judged acceptable because it appeared acceptable from that perspective—and not because it's judged true, from the more informed evaluator's perspective, that the protagonists actually knew something false.

There are now at least two hypotheses available to explain linguistic evidence like Know Ulcers. One takes such sentences as literal knowledge ascriptions that challenge Factivity. The other takes such sentences as non-literal ascriptions and instances of protagonist projection that do not challenge Factivity. Five experiments were conducted to test these two hypotheses.

3   Experiments in Projection

The following experiments test whether protagonist projection can account for seemingly non-factive use of certain purportedly factive verbs in ordinary language. The key question is whether or not people really do literally attribute knowledge to subjects with false beliefs, or if such statements are better understood in terms of protagonist projection.

3.1   Experiment 1

This experiment tests whether protagonist projection is active in Know Ulcers, Know Sailing, Learn History, and Realize Death.[7] Participants ($N = 73$, 37 female, median age = 35) were randomly presented with items very similar to each of the four cases considered as counterevidence to Factivity in Section 1. These included two sentences including the verb 'know' and one sentence each involving 'learn' and 'realize'.[8] For example, here is one of the four items given to each participant:

> **Everyone knew** that stress caused ulcers, before two Australian doctors in the
>
> early 1980s proved that ulcers are actually caused by bacterial infection.

Directly after seeing each of the items, participants were given questions of the following form, respective to the different verbs imbedded. They were then asked to select one of the explicit paraphrases provided regarding the usage of that verb. The order of these

---

[7] All studies used the internet-based commercial research tools Mturk and Qualtrics. Online samples were restricted to participants located in the United States. Participants were not allowed to re-take any study reported here. Upon completion participants saw a short demographic questionnaire. Ninety-eight percent of all participants reported English as a native language. Preliminary analysis revealed no effects of participant age, gender or native language for each of the experiments presented in the paper, so the analyses collapse across those demographic variables.

[8] A list of materials is available for Experiment 1 and Experiment 2 in the appendix.

items and answer choices was randomized.[9] Below is the question that was paired with the ulcer case above:

> Which of the following do you think best **describes** what is meant in the **BOLD portion** of the above sentence:
>
> A) Everyone thought they knew
>
> B) Everyone really did know

If participants are reading these sentences and interpreting the verbs as factive, then we would expect them to be more likely to adopt the projective answer—that from protagonist x's perspective, 'x thought that x Φ'. Alternatively, if the items really are best interpreted as counterexamples to Factivity, then we would expect participants to give the non-projective answer—that from the participant's more informed perspective, "x really did Φ."

What was found was a split in the way participants answered for 'knows' on the one hand, and 'learns' and 'realizes' on the other. Participants were significantly more likely to give the non-projective "really Φ" answer for both of the sentences containing 'learns' and 'realizes' than for both of the sentences containing 'knows', McNemar's test, $ps < .01$. Only 9% in the ulcer case and 22% in the sailing case gave the non-projective answers to the knowledge sentences, which both fell far below chance rates, binomial test, $ps < .01$, test proportion = .5, while 56% in the history case and 49% in the death

---

[9] This explicit paraphrase task has also been used successfully in other domains, including the study of mental state attributions to groups (see Phelan 2010, Phelan et al. 2013).

case gave the "really Φ" non-projective answer to the learning and realizing sentences, which both did not differ from chance rates, binomial test, $p = .349/p > .99$, test proportion = .5. These results are shown in Figure 1.
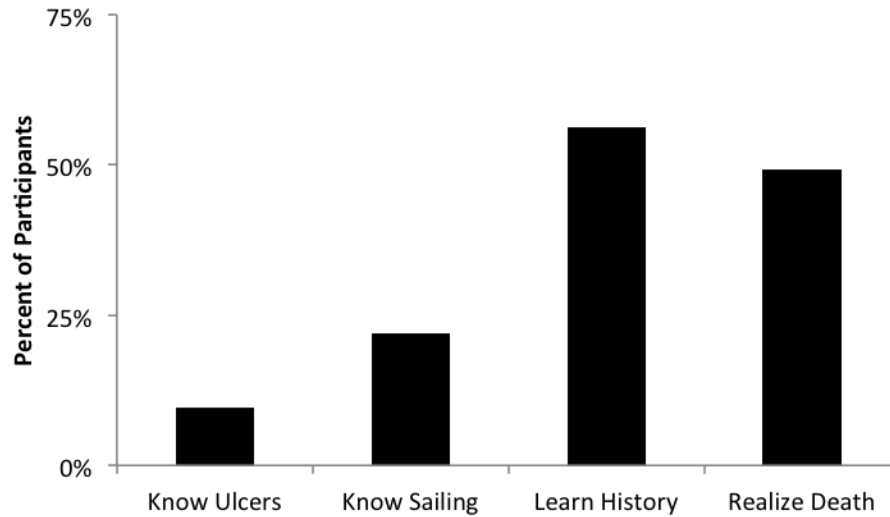

Figure 1: Experiment 1. Percent of participants giving "really Φ" answers.

Interestingly, there seems to be some disagreement about the best interpretation of the sentences involving 'learns' and 'realizes'.[10] Perhaps this is only a result of the particular sentences selected. Or perhaps higher non-projective scores indicate that there may be literal non-factive readings of 'learns' and 'realizes'. By comparison however, there seems to be little disagreement about the best interpretation of the sentences involving 'knows'. The data indicate that participants are significantly more likely to interpret the

---

[10] Another explanation is that a salient reading of 'learn that' is 'is taught that', and of course 'is taught that' certainly isn't factive. Many thanks to John Turri for discussion on this point.

sentences involving 'knows' from the protagonists' perspective rather than as literal ascriptions of knowledge.

3.2   Experiment 2

Findings in Experiment 1 give us reason doubt that the sentences considered in Section 1 count against Factivity. But it is still a bit unclear what these data are telling us. After all, the cases differed not only in the verb used (e.g. knows vs. realized) but also in the various complement clauses expressing the false propositions. Experiment 2 sought to replicate projective findings for 'knows' in Experiment 1, but this time by displaying the different factive verbs in the same contexts and with the same complement clauses. Also, the verb 'believe' was included in this study, serving as an additional experimental comparison for the other verb findings.

In a between-subjects design, participants ($N = 120$, 70 female, median age = 34) were presented with a vignette about two biologists trying to identify a crab. While each participant only saw one of four possible stories, each story varied by which verb was embedded: 'know', 'realize', 'learn', or 'believe' (a full list of materials appears in the appendix). For example, below are the materials that participants saw involving knowledge:

Two biologists are having a discussion about a crab:

Dr. Krycek: "Can we be sure that this one is of the genus *Calcinus*?"

Dr. Scully: "I have my notes right here.
I know that this is a specimen of *Calcinus hazletti*"

But this crab was actually a completely different crab, *Calcinus obscurus.*

After each story, participants were asked the same basic question involving the relevant verb of the vignette they were given. Below is the question seen by participants given the knowledge case above:

Which of the following do you think best **describes** what is meant in the **BOLD portion** of the above sentence:

1. Dr. Scully thought she knew

2. Dr. Scully really did know

As in Experiment 1, if participants think that Dr. Scully's utterance, 'I know that this is a specimen of *Calcinus hazletti*' actually describes a true state of affairs, then we would expect them to answer from their perspective that 'Dr. Scully really did know'. Alternatively, if participants interpret this statement by projecting, then the prediction would be that they would be more likely to answer that 'Dr. Scully thought she knew'. In other words, this answer would count as good evidence that the reason participants find the incorrect thing Dr. Scully said acceptable is because it is true from the protagonist's perspective that she thought she knew, and not because she knew something false.

There was again a significant split in answers between verbs. Only a very small percentage of participants gave the "really Φ" non-projective answer in the knowledge condition compared to the belief condition, 10%/63%, Fisher's exact test, p < .01. Moreover, answers in the knowledge condition fell far below chance rates, binomial test, $p < .01$, test proportion = .5. People were once again more likely to give "really Φ" non-projective answers for both 'learn' and 'realize', 33%/36%, than for 'knows', Fisher's exact test, $ps < .05$, but were significantly less likely to give the "really Φ" answer for both 'learn' and 'realize', than for 'believe', Fisher's exact test, $ps < .05$. Rates of non-

projective answers were trending from chance in the learn condition, binomial test, $p =$ .099, test proportion $= .5$, but did not differ significantly from chance in either the realize or belief conditions, $p = .15/.20$. These results are displayed in Figure 2.
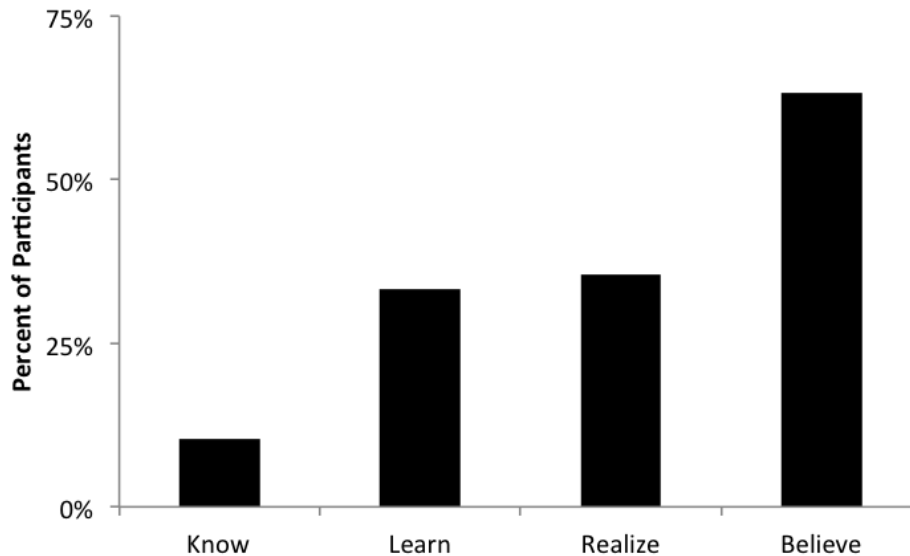

Figure 2: Experiment 2. Percent of participants giving "really Φ" answers.

The results of Experiment 2 replicate the main finding of Experiment 1. The results again question the claim that protagonists with false beliefs are viewed as having knowledge. Compared to the other verbs included across Experiment 2, participants were significantly more likely to give projective readings of 'knows'. It seems unlikely then that these knowledge sentences are best interpreted as counterexamples to Factivity rather than as non-literal ascriptions of knowledge.

3.3 Experiment 3

Experiment 2 continues to build a case in support of Factivity. But although it showed that people are more likely to give projective rather than non-projective readings for

'knows' than 'believes' in cases with identical false complement clauses, this does not give us evidence that the false complement clauses are directly responsible for their answers. It's possible that participants only tended to give projective answers to the sentences because they did not think the subjects had actually Φ-ed, independently of whether or not they think the verbs are or aren't factive. For example, participants may have answered that Dr. Scully only 'thought she learned' about the crab because she made an error when copying down her notes, or misremembered the lesson, and not because it's impossible to learn false things.

Experiment 3 sought to find evidence that the truth or falsity of the proposition in question is what is directly responsible for projective or non-factive answers across the various verbs tested. In a between-subjects design, participants ($N = 217$, 115 female, median age = 34) were presented with another vignette adapted from Hazlett involving two police offers relaying some testimony. While each participant only saw one story, the stories varied only by the verb used in bold ('know', 'learned', 'realized', or 'believe'). Below is the vignette involving 'knows':

> Officer Ted asks the Police Sergeant, "Is there any information from the FBI about how the bomb was constructed?" The Sergeant told him, "From the investigation, **they know** the bomb was homemade."

After seeing one vignette, roughly half of participants were presented with a complement phrase expressing a false proposition.

> *False Complement*: But actually, the bomb in question only appeared homemade. Instead, professionals constructed it in a high-tech chemical plant.

The other half received a complement expressing a true proposition:

True Complement: And as it turns out, the bomb in question was homemade. Non-professionals had constructed it in a basement apartment.

Lastly, participants were presented with the same question used in Experiments 1-3 to test for projective interpretations of the various different verbs:

Which of the following do you think best **describes** what is meant in the **BOLD portion** of the above sentence:

A) The FBI thought they knew

B) The FBI really did know

If the results support Factivity, then we would expect a large majority of participants presented with the factive verbs to answer that the 'FBI really did Φ' when the accompanying complement clause expresses a true proposition, and that 'FBI thought they knew' when given a false complement. Again, projection would predict that the incorrect testimony relayed by the sergeant when she said 'they know' is acceptable to participants only in so far as it is true from the protagonists' perspective, and not because false things are known. Alternatively, we would expect no significant effect of truth or falsehood on projective responses in belief conditions.

That was exactly what was found. Participants gave the non-projective, 'really Φ' answer for 'know' (86%), 'learn' (78%), 'realize' (91%), and 'believe' (80%) in each of the true complement conditions at rates that significantly exceeded chance, binomial test, *ps* < .05, test proportion = .5. By comparison, participants were significantly less likely to give the 'really Φ' answer for 'know' (12%), 'learn' (21%), 'realize' (31%), in each of the false compliment conditions than they were in each of the true compliment conditions, Fisher's exact test, *ps* < .01. Scores in each of these conditions were

significantly below chance rates, binomial test, ps < .05, test proportion = .5. However assignment to true or false complement condition had no significant effect on judgments about 'believe', 80%/61%, Fisher's exact test, *p* = .23, and scores in the false belief condition did not differ from chance, binomial test, *p* = .24, test proportion = .5. Lastly, answers for 'know', 'learn', and 'realize' were each significantly lower than 'believes' in false compliment conditions, Fisher's exact test, *ps* < .05. These results are shown in Figure 3.
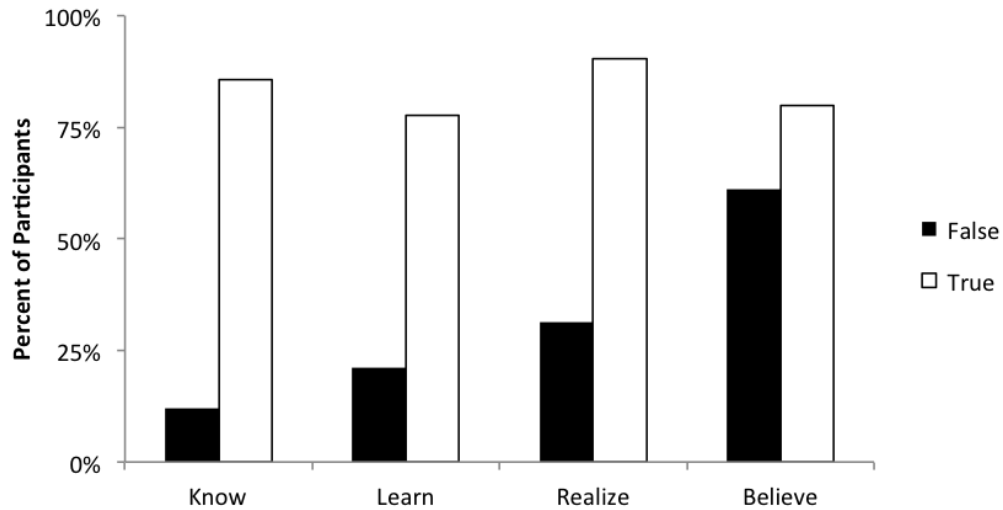

Figure 3: Experiment 3. Percent of participants giving "really Φ" answers grouped by true or false complement.

3.4 Experiment 4

The pattern of results observed across Experiments 1-3 point toward Factivity for 'knows'. However these experiments used the same basic knowledge probe. To ensure that responses were not an artifact of this manner of questioning, the subsequent follow-up experiment was conducted.

Participants ($N = 63$, 32 female, median age = 33) saw the same basic stories and experimental design used in Experiment 3 involving officers relaying testimony. Participants were presented with one of four possible cases that varied the truth or falsity of the complement clause, as well as the verb imbedded (either 'knows' or 'believes'). Below is an example of the false complement materials varying these two verbs:

Officer Ted asks the Police Sergeant, "Is there any information from your contact at the FBI about how the bomb was constructed?" The Sergeant told him, "Yes. Mr. Smith told me that, from his investigation, he [knows/believes] the bomb was homemade." But actually, the bomb in question only appeared homemade. Instead, professionals constructed it in a high-tech chemical plant.

Directly after these stories participants then saw the following questions:

According the story, was the bomb actually homemade? [Yes/No]

Which do you think best describes Mr. Smith?

A. Mr. Smith only thought he [knew/believed] the bomb was homemade

B. Mr. Smith [knew/believed] the bomb was homemade

Note that these probes varied in a number of ways from the probes used in the previous experiments. First, the cases involved a particular FBI agent named Smith. Given work on group attributions, one might suspect that asking about a particular person rather than a group might promote more strict literal readings (Phelan et al. 2013). Second, participants were asked, 'Which do you think best describes Agent Smith?' just in case the particular previous wording of the question or method of bolding of verbs was contributing to different readings of 'know'. Third, answer choices dropped the intensifier (now 'Agent Smith only thought he Φ the bomb was homemade' or 'Agent

Smith Φ that the bomb was homemade'). Lastly, unlike Experiment 3 these participants were given a comprehension check explicitly acknowledging the truth or falsity of the complement.

The result was a powerful replication of Experiment 3.[11] Only 7% of participants in the false knowledge condition gave the non-projective answer, which was far below chance rates, binomial test, $p < .01$, test proportion $= .5$. Responses in the false knowledge condition differed significantly from responses to each of the other three conditions, Fisher's exact test, $ps < .01$, with the overwhelming majority giving the non-projective "Agent Smith Φ-ed" answer in the false belief, true belief and true knowledge conditions, 93%/94%/83%, which each far exceeded chance rates, binomial test, $ps < .01$, test proportion $= .5$. These results are shown in Figure 4.
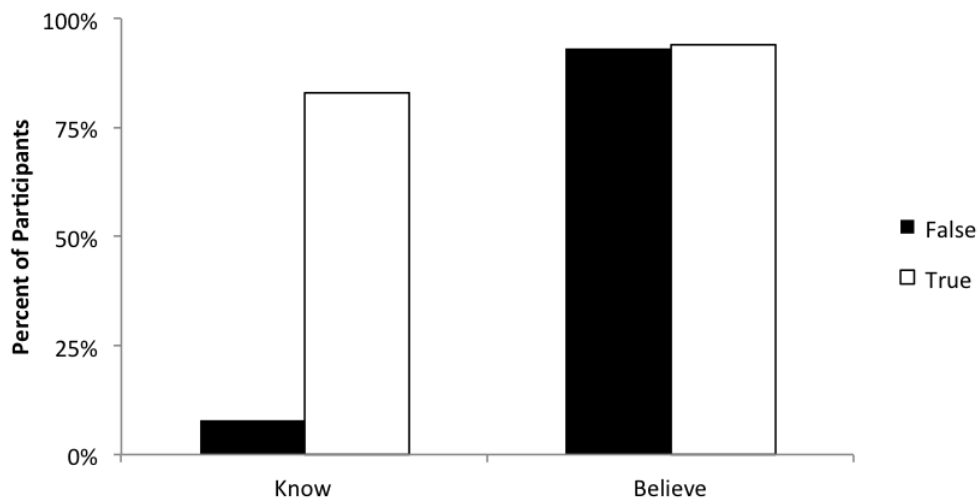

Figure 4: Experiment 4. Percent of participants giving the "Agent Smith Φ-ed" answer grouped by true or false complement.

---

[11] Seven participants were removed for failing the comprehension check.

3.5   Experiment 5

The results from Experiments 1-4 strongly support Factivity for 'knows'. They have done so by taking cases and materials previously presented as evidence against Factivity and demonstrating that they are readily explained by protagonist projection. Turning from prior challenges however, we might wonder about the phenomenon of protagonist projection itself. The experiments thus far have shown a strong tendency for protagonist projection when participants are asked about the uses of certain verbs in short sentences. Does protagonist projection also occur when participants are asked to consider longer vignettes and make their own judgments about them?

Experiment 5 sought to test this possibility. Given the extensive research done in traditional and experimental epistemology on what are known as "bank cases" to date (DeRose 1992, Buckwalter 2010, Schaffer & Knobe 2012), these cases were also chosen to further investigate projective ascriptions. Since these cases have typically been used to study the salience of attributor error (or the possibility that a mistake might be made) they seem like apt cases to investigate the effects that false beliefs (or that an actual mistake has been made) may have on ordinary knowledge judgments. Rather than supplying participants with non-factive usage, the idea is to ask participants to make knowledge judgments about protagonists in these cases and then test the actual judgments people make for occurrences of projection.

In a between-subjects study, all participants in Experiment 5 ($N = 126$, 77 female, median age = 34) saw the following text, very similar to the original bank cases, and to previous research done with bank stimuli (see Buckwalter 2010):

Hannah and her sister Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. As they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Since they do not have an impending bill coming due, and have plenty of money in their accounts, it is not important that they deposit their paychecks by Saturday. Hannah says, "I was just at this bank two weeks ago on a Saturday morning, I know that the bank will be open on Saturday. Let's leave and deposit our paychecks tomorrow morning."

After reading this basic story, participants randomly saw one of two different conclusions to the story varying whether the bank was actually open on Saturday was either true or false:

True belief: When they returned the next day, they found that the bank was open for business on Saturday, as it usually is most mornings.

False belief: When they returned the next day, they found that the bank was closed on Saturday due to a lightning strike that occurred earlier that morning.

Each participant was then asked three questions. The first question was designed to serve as a basic comprehension and manipulation check to insure that participants read the story, as well as understood that the subjects above had explicitly true or false beliefs:

According to the story, was the bank actually open for business on Saturday morning? [Yes/No]

Consistent with past research on these cases, on a five-item scale (1 being "definitely false," 3 being "neither true nor false," and 5 being "definitely true") participants were asked to judge the knowledge statement made by the subject in the story above:

When Hannah said, "I know that the bank will be open on Saturday," is what she said true or false?

Lastly, participants were then asked the question designed to reveal whether or not protagonist projection was present:

Which do you think best describes Hannah in the story?

A. Hannah thought she knew

B. Hannah really did know

The result of the study was that in accord with traditional philosophical wisdom, the truth of the epistemic subject's belief had a large effect on participant's knowledge judgments. Participants were less likely to agree that the protagonist's knowledge claim was true when the protagonist's belief was false ($M = 3.71$, $SD = 1.34$) then when it was true ($M = 4.52$, $SD = 0.57$), $t(104) = -4.18$, $p < .01$.[12]

The impact of truth on participant knowledge judgments lends additional support to Factivity. At the same time, mean answers in the false belief condition were nonetheless still significantly above the neutral midpoint of three on this five item scale, $t(47) = 3.67$, $p < .01$, suggesting that some participants may consider Hannah's knowledge sentence 'I know that the bank will be open on Saturday,' to be true when her belief is false. In this case, the further question now remains as to whether or not this is evidence of an instance of non-factive knowledge or for projection. To find out we need to look at people's answers to the projection question. Again, for judgments that Hannah's knowledge sentence is true to count as good evidence against Factivity, we would expect

---

[12] Twenty participants were removed for failing comprehension check.

them to be just as likely to answer 'Hannah really did know' despite the truth or falsity of her belief. Alternatively, if people are projecting, then we would predict significantly more "Hannah really did know" answers when Hannah's belief is true, and more "Hannah thought she knew" answers when Hannah's belief is false.

Assignment to condition did affect the rates at which participants selected projective answers. Sixty-nine percent of participants answered that Hannah really did know' in the true belief condition, whereas 60% answered that 'Hannah thought she knew' in the false belief condition, Fisher's exact test, $p < .01$. Answers in the true belief condition differed from chance rates, binomial test, $p < .01$, test proportion = .5, but answers in the false belief condition did not reach statistical significance from chance, $p = .193$.

These results show that some participants do appear willing to answer that 'S knows that p' when p is false. But participants are also more prone to answer from the perspective of the protagonist, rather than literally attribute knowledge to the protagonists, when the protagonist's belief is false. These findings lend additional support to the claim that at least one major conception of knowledge coheres with the factive account. At the same time, it is noteworthy that projective scores in the false condition are at chance rates, and lower than those observed in Experiments 1-4.

A number of things might explain why projective scores were lower in this particular study. For one, this could be the result of an order effect. The projection instrument always followed a knowledge ascription question, the answers to which were significantly affected by truth and might have interfered with these responses. For another, the test item "When Hannah said, "I know that the bank will be open on

Saturday,' is what she said true or false" could have very well been interpreted as, "whether what she thought she knew was true or false," just as we would predict if ascription was prone to implicit protagonist projection.[13] Lastly, the materials in Experiment 5 involve knowledge claims of future events, which might also differentially affect rates of projection. Future work might profitably explore these possibilities and thereby expand our understanding of the frequency and limits of protagonist projection.

4   Support for Factivity

We began with a discussion of the standard truth condition in epistemology, and whether or not patterns of ordinary use are completely at odds with a factive account. We then considered several sentences that seem to be acceptable or non-deviant to speakers, but also appear to use factive verbs—including 'knows'—in transparently non-factive ways. Two hypotheses were presented to explain these linguistic data. The first was that they constitute literal ascriptions of knowledge under false belief. The second was that these ascriptions are non-literal, and thereby fail to challenge ordinary support for the truth requirement on knowledge.

The results of Experiments 1-5 point to the important and systematic role of truth in knowledge judgments. Such evidence does not rule out the existence of a non-factive concept or show that knowledge ordinarily understood is necessarily factive. But it does show that one can explain many apparent uses of factive verbs in cases where complement clauses express false propositions without needing to call into question

---

[13] Many thanks to an anonymous reviewer for discussion on these points.

widespread ordinary language support for the factivity of 'knows'. One mechanism for doing so is protagonist projection.

In Experiment 1 we saw that people were willing to adopt the perspective of the protagonists in cases considered previously as evidence against factivity. In Experiments 2-3, we saw that projective paraphrasing of factive verbs arose because of the truth or falsity of the complement clause across closely matched contexts of usage. Experiment 4 replicates and expands these findings using different question options. Lastly, Experiment 5 demonstrates that projection may even occur when participants are presented with the kinds of epistemic thought experiments often considered by philosophers and asked to make their own judgments about them.

There were also two interesting and unexpected outcomes of these experiments. The first was that there may be differences between ordinary uses of 'knows' and other purportedly factive verbs, e.g. 'learns', and 'realizes'. Participants in Experiment 1-2 gave significantly more non-projective answers for these verbs than for 'know'. This effect in Experiment 3 did not reach statistical significance, but was suggestive in the same direction. More research is needed to discover whether literal non-factive readings really are available for these verbs or if these effects are artifacts of present experimental materials. But at the very least, it is beginning to seem like there is something special about 'knows' that doesn't also show up as readily for these other verbs. By comparison, people seem to be adamantly adopting projective paraphrases rather than genuinely ascribing knowledge to protagonists with false beliefs in the cases considered.

The second interesting and unexpected outcome was that there may be subtle variation or individual differences in knowledge judgments relating to factivity.

Experiments 1-4 exhibited surprisingly consistent and powerful results in favor of Factivity (with only 7%-22% answering otherwise). Nevertheless, further research might profitably investigate why this small minority was willing to make non-projective ascriptions. Additionally, while there was an effect for projection between conditions in Experiment 5, a higher percentage of participants gave the non-projective answer than in each of the other experiments (with up to 40% answering otherwise). Comparing results between experiments is fraught. For present purposes, the evidence in Experiment 5 demonstrates that truth has a significant influence on both explicit knowledge judgments and projective reinterpretations, and thus, continues to point toward Factivity. Nonetheless this does not eliminate the possibilities that a substantive minority exists that are either continually willing to literally ascribe knowledge to falsehoods or non-literally ascribe knowledge in a way that is not explained by projection.

Before moving on it is worth noting that the former possibility is unlikely. Many studies on knowledge attributions have found evidence that truth is relevant to knowledge judgments among adults as well as very young children (Shatz et al 1983; Sodian 1988; Booth et al. 1997; Turri, Buckwalter & Blou in press; for additional work including mediation studies see Buckwalter & Turri under review). Orthodoxy has stood for centuries and detractors to Factivity have yet to provide empirical evidence to question it among ordinary speakers. At the same time, one limitation of the present studies is that the question of ordinary language support for the truth requirement is found only in terms of the English word 'know'. It is an open possibility that variation in language, culture or other individual differences yet to be studied would yield a different verdict, and no arguments or data presented here forestalls this possibility.

The evidence suggests that philosophical orthodoxy does track one major conception of knowledge predominately used by ordinary people. This undercuts current challenges to Factivity, and vindicates ordinary langue support for the standard truth condition. Of course, neither these nor any data could show that all putative non-factive uses of 'know' are cases of protagonist projection. Nor do these data show that there cannot be evidence against Factivity independently of projection phenomena. What the data do suggest is that dissidents must collect more empirical evidence before we should seriously consider abandoning orthodoxy on the basis of ordinary language ascriptions.

Far from a radical departure from the things philosophers have said about knowledge, the results lend support for the standard truth condition. For all its appeal, there has actually been surprisingly little by way of explicit argumentation for the standard truth condition. Often it is simply assumed or stipulated in philosophical analysis (see Schaffer 2005 for one example). Just as ordinary ascription practices have been used by some as key evidence for the relationship between knowledge and practical interests or the propriety of deciding on courses of action (Hawthorne 2004; Stanley 2005; Hawthorne and Stanley 2008), so too can ordinary practice add empirical warrant to the orthodox claim that knowledge is composed partially of truth.

The systematic effect that the truth or falsity of a belief has on ordinary ascriptions also strongly suggests that truth is an important component of the ordinary concept of knowledge. This claim fits well with observations made in prior research about the standards governing instructional demonstration and assertion (Buckwalter and Turri 2014; Turri 2013). This work on the give and take of ordinary conversational exchanges demonstrates that knowledge has a large impact on whether people think an

assertion should be made, and that this requires that the assertion in question is true. Taken together, this is more evidence that the concept of knowledge that people tacitly hold is factive (though see the cautionary note above).

Lastly the evidence inspires methodological caution when appealing to empirically unexamined ordinary language practices in epistemology. Many philosophers rely on evidence from ordinary usage (for variations on this theme see Cohen 1988; Lewis 1996; Schaffer 2005; Stanley 2005). In fact, accounting for ordinary "knowledge-citations" (Fantl and McGrath 2007, p. 562), "ordinary folk appraisals of the behavior of others" (Hawthorne and Stanley 2008, p. 571) and "our inclinations to ascribe knowledge" (Nagel 2008, p. 279) has largely contributed to the discovery and development of new leading theories of knowledge (DeRose 2009; Hawthorne 2004; Stanley 2005). The present studies aim to contribute to this same basic approach. But the studies also suggest that ordinary epistemic behaviors can be quite subtle. Epistemologists and cognitive scientists should be very careful to examine these subtleties before drawing sweeping conclusions about their import on theories of knowledge. This paper has provided one procedure for doing so, namely for distinguishing between literal and non-literal projective ascriptions of knowledge.

5   Conclusion

We know that we are surrounded by non-factive uses of factive verbs like 'knows'—or at least—we thought we knew it. Under controlled experimental conditions in which participants are given explicit paraphrasing tasks, a more likely explanation of the linguistic evidence considered is that participants are adopting projective readings, rather

than literally attributing false knowledge. These projective readings undercut existing arguments against Factivity and point to the powerful and important relationship between knowledge and truth in ordinary judgments of knowledge.[14]

REFERENCES

Armstrong, D. (1969). Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70, 21–36.

Austin, J. L. (1956). A plea for excuses. Proceedings of the Aristotelian society, 57, 1–30.

Beebe, J. R. (2014). *Advances in experimental epistemology*. Bloomsbury Academic.

Booth, J. R., Hall, W. S., Robison, G. C., & Kim, S. Y. (1997). Acquisition of the mental state verb know by 2- to 5-year-old children. *Journal of Psycholinguistic Research*, 26(6), 581–603. doi:10.1023/A:1025093906884

Buckwalter, W. (2010). Knowledge Isn't Closed on Saturdays. *Review of Philosophy and Psychology* 1: 395-406.

Buckwalter, W. (2012). Non-traditional factors in judgments about knowledge. *Philosophy Compass* 7/4 (2012): 278–289.

Buckwalter, W. and Turri, J. 2014. Telling, Showing and Knowing: A Unified Account of Pedagogical Norms. *Analysis* 74 (1): 16-20.

Buckwalter, W., Rose, D., and Turri, J. (in press). Belief through thick and thin. *Nous*.

Buckwalter, W., and Turri, J. (under review). Descartes's Schism, Locke's Reunion: Completing the Pragmatic Turn in Epistemology.

Chisholm, R. (1989). *Theory of Knowledge*, 3rd ed. Prentice-Hall.

Cohen, J. (1966). More about Knowing and Feeling Sure. *Analysis*, 27, 11–16.

Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives* 2: 91-123.

Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13: 57–89.

DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52: 913-29.

DeRose, K. (2009). *The case for Contextualism,* Oxford, Oxford University Press.

Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, Mass.: MIT Press.

Fantl, J. & M., McGrath (2009). *Knowledge in an Uncertain World*. Oxford University Press.

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

Friedman, O., & Turri, J. (in press). Is probabilistic evidence a source of knowledge? *Cognitive Science*.

Goldman, A. I. (1986). *Epistemology and Cognition*. Harvard University Press.

Hawthorne, J. (2004). *Knowledge and Lotteries*. Clarendon Press.

Hawthorne, J. and J., Stanley (2008). Knowledge and Action. *Journal of Philosophy* 105: 571–590.

Hazlett, A. (2010). The myth of factive verbs. *Philosophy and Phenomenological Research* 80: 497-522.

Hazlett, A. (2012). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica* Online First: http://dx.doi.org/10.1007/s12136-012-0163-3.

Holton, R. (1997). Some Telling Examples: A reply to Tsohatzidis. *Journal of Pragmatics* 28: 624-28.

Jones, O. R. (1971). Knowing and Guessing: By Examples. *Analysis*, 32, 19–23.

Kiparsky & Kiparsky (1971). Fact. Reprinted in *Semantics*. Steinberg, D., and Jakobovitz, L. Cambridge: Cambridge University Press, 345-369.

Lewis, D. (1996). Elusive knowledge. Reprinted in *Papers in Metaphysics and Epistemology,* Cambridge: Cambridge University Press, pp. 418-445.

Lehrer, K. (1974). *Knowledge*. Clarendon Press.

Marshall B., J., Armstrong, J. A., McGechie, D.,B., & Glancy, R.J. (1985). Attempt to fulfill Koch's postulates for pyloric campylobacter. *Medical Journal of Australia* 142: 436-439.

Moore, G.E. (1942). A Reply to My Critics. In *The Philosophy of G.E. Moore* (La Salle, Ill.: Open Court), ed. P. Schilpp.

Murray, D., Sytsma, J. and Livengood, J. 2013. 'God Knows (but does God Believe?).' *Philosophical Studies*, 166: 83-107.

Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Nous*, 47(2), 371–384.

Nagel, J. (2008). Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy* 86: 279–294.

Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. Cognition, 1–10. doi:10.1016/j.cognition.2013.02.008

Nolan, D. (2008). Non-Factivity About Knowledge: A Defensive Move. *The Reasoner* 2.11: 6-7.

Nozick, R. (1981). *Philosophical explanations*. Cambridge, Mass: Harvard University Press.

Phelan, M. (2010). The inadequacy of paraphrase is the dogma of metaphor. *Pacific Philosophical Quarterly* 91: 481-506.

Phelan, M., Arico, A. and Nichols, S. 2013. Thinking Things and Feeling Things: On an Alleged Discontinuity in Folk Metaphysics of Mind. *Phenomenology and the Cognitive Sciences*, 12: 703-725.

Plato, ., Cooper, J. M., & Hutchinson, D. S. (1997). *Complete works*. Indianapolis, Ind: Hackett Pub.

Reid, T. (1785/2002). Essays on the intellectual powers of man. (D. Brookes, Ed.) University Park: Pennsylvania State University Press.

Rose, D. and Schaffer, J. 2013. Knowledge Entails Dispositional Belief. *Philosophical Studies*, 166:19-50.

Ross, W.D., (1984). *The Complete Works of Aristotle*. The Revised Oxford Translation, vol. 2, Jonathan Barnes, ed., Princeton: Princeton University Press.

Russell, B. (1948/2012). *Human knowledge: its scope and limits*. New York: Routledge.

Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Nous* 35: 477–514.

Sackris, D., and J., Beebe. (2014) Is Justification Necessary for Knowledge? In Beebe, J. R. (Ed.). *Advances in experimental epistemology*. Bloomsbury Academic.

Schaffer, J. (2005). Contrastive Knowledge. *Oxford Studies in Epistemology* 1: 235-71.

Schaffer, J. and Knobe, K. 2012. Contrastivism Surveyed. *Noûs*, 46: 675-708.

Shatz, M., Wellman, H. M., & Silber, S. (1983). The acquisition of mental verbs: A systematic investigation of the first reference to mental state. *Cognition*, 14(3), 301–321.

Sellars, W. (1963). *Science, perception and reality*. Atascadero, CA: Ridgeview Publishing Company.

Sodian, B. (1988). Children's attributions of knowledge to the listener in a referential communication task. *Child development*, 59(2), 378–385.

Sosa, E. (1991). Knowledge in Perspective. Cambridge: Cambridge University Press.

Stanley, J. (2005). *Knowledge and Practical Interests*, Oxford University Press.

Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford: Clarendon Press.

Starmans, C. and Friedman, O. 2012. The Folk Conception of Knowledge. *Cognition,* 124: 272–283.

Steup, M. (2006). The Analysis of Knowledge. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/knowledge-analysis/#IEX>.

Tsohatzidis, S. (2012). How to Forget that "Know" is Factive. *Acta Analytica* DOI: 10.1007/s12136-012-0150-8.

Turri, J. (2011a). Manifest Failure: The Gettier Problem Solved. *Philosophers' Imprint* 11.8, 1-11.

Turri, J. (2011b). Mythology of the Factive. *Logos & Episteme* 2: 143-52.

Turri, J. (2013). The test of truth: an experimental investigation of the norm of assertion. *Cognition* 129: 279–91.

Turri, J. Knowledge as Achievement, More or Less. Forthcoming in *The Present and Future of Virtue Epistemology*, ed. Miguel Angel Fernandez, press TBD.

Turri, J., Buckwalter, W., & Blouw, P. (in press). Knowledge and Luck, *Psychonomic Bulletin & Review*.

Vogel, J. (1990). Are there Counterexamples to the Closure Principle In M. D. Roth & G. Ross (Eds.), *Philosophical Studies Series* (Vol. 48, pp. 13–27). Springer.