Episteme, 15, 4 (2018) 383–411 © Cambridge University Press
doi:10.1017/epi.2017.7

COULD KNOWLEDGE-TALK BE LARGELY NON-LITERAL?

JULIANNE CHUNG
julianne.chung@louisville.edu

ABSTRACT

Much work in epistemology, old and new, is devoted to explaining – or explaining away – the appeal of skeptical arguments (that is, arguments designed to impugn most, if not all, of our claims to know), as well as the fact that it also often seems as if we use knowledge-attributing sentences to express truths. More recently, philosophers have sought to explain other kinds of variability in our use of knowledge-attributing sentences besides, such as the kinds of variability that occur as a result of changing what is at stake or the salience of alternative possibilities, as in scenarios such as Keith DeRose's "bank cases." These cases, and others like them, pose a dilemma very much like that posed by skeptical arguments. For they appear to suggest that we are wrong about something. Either we speak falsely when we attribute knowledge in the situations under consideration, or we speak falsely when we deny it. Or so it would seem. This paper argues that an alternative approach to explaining the relevant variability in our use of knowledge-attributing sentences – namely, a fictionalist one – merits serious consideration.

INTRODUCTION

Much work in epistemology, old and new, is devoted to explaining – or explaining away – the appeal of skeptical arguments (that is, arguments designed to impugn most, if not all, of our claims to know), as well as the fact that it also often seems as if we use knowledge-attributing sentences to express truths.[1] More recently, philosophers have sought to explain other kinds of variability in our use of knowledge-attributing sentences besides, such as the kinds of variability that occur as a result of changing what is at stake or the salience of alternative possibilities, as in scenarios such as Keith DeRose's "bank cases."[2] These cases, and others like them, pose a dilemma very much like that posed

1 And knowledge-denying sentences to express falsehoods. In what follows I generally only speak of knowledge-attributing sentences; however, this restriction would have to be removed in a final analysis.

2 For DeRose's statement of these cases, as well as the problems that they present, see DeRose (2009: 1–2). DeRose uses these cases to motivate the claim that would-be knowledge-attributors are sensitive to what is at stake for agents that they are considering attributing knowledge to, as he expects that many will share his intuition that while knowledge attributions made in low-stakes scenarios often seem true, similar knowledge attributions made in high-stakes scenarios often seem false, even when the relevant agents have the same evidence. According to DeRose, if an agent, say, wants to deposit a check on a particular Saturday, we are normally inclined to think it true that the agent knows that the bank will be open then if the agent recalls being there on that day two weeks prior – but not if (for example)

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

by skeptical arguments. For they appear to suggest that we are wrong about something. Either we speak falsely when we attribute knowledge in the situations under consideration, or we speak falsely when we deny it. Or so it would seem. Because of this, proposed ways of explaining the relevant variability in our use of knowledge-attributing sentences are very similar as well. Five approaches in particular have proved highly influential in contemporary Western epistemology. Here are the central tenets of each:

(i) We are wrong to attribute knowledge so generously (but there is a good explanation for why we do).[3]
(ii) We are wrong to abandon knowledge attributions at the first sign of doubt (but there is a good explanation for why we do).[4]
(iii) Semantic features of sentences that (at least appear to) attribute knowledge are more complex than one might have first thought, as the best semantics for knowledge-attributing sentences is, e.g., contextualist or relativist.[5]
(iv) The metaphysics of knowledge is more complex than one might have thought at first, as, e.g., subject-sensitive or interest-relative invariantism with respect to knowledge is true.[6]
(v) Pragmatic features of our use of sentences that attribute knowledge are more complex than one might have first thought, such that we are often warranted in asserting knowledge-attributing (or -denying) propositions despite their (potential) falsity.[7]

---

the agent has just written a very important check, a check that will only clear if the deposit is successfully made on that Saturday. Similar cases can be found (for example) in the work of others who defend semantic explanations of variability in our use of knowledge-attributing sentences, as well as in the work of defenders of views such as subject-sensitive or interest-relative invariantism.

3 These are often called "skeptical strategies"; for a discussion of such approaches, see, e.g., Unger (1971, 1975), Stroud (1984), Nagel (1986), and Hawthorne (2004).

4 These are often called "anti-skeptical strategies." Such approaches have proved the most popular in contemporary Western epistemology, at least with respect to the project of explaining skeptical debates if not scenarios like DeRose's bank cases. For this reason, I will not be able to provide anything close to a comprehensive list of such attempts here. With that, (just) a few relatively recent examples can arguably be found in: Dretske (1970), Goldman (1976, 1979), Stine (1976), Klein (1981, 1995), Nozick (1981), Putnam (1981), Vogel (1990a), Sosa (1994), Lehrer (2000), Pryor (2000), Williamson (2000), and Wright (2004).

5 Examples of this approach include contextualist (see, e.g., DeRose 1995, 2009; Lewis 1996; Cohen 1999, 2000a), and Blome-Tillman (2009, 2014), contrastivist or comparativist (Schaffer 2004a; Schaffer and Szábo 2014) and relativist (MacFarlane 2005, 2011) accounts of knowledge attributions. I say, "at least appear to" here because some theories, such as expressivist ones (which are gaining traction), may maintain that such sentences do not in fact attribute knowledge.

6 See, for example, Hawthorne (2004), Stanley (2005), and Fantl and McGrath (2009).

7 Approaches that fit this description fall under the banner of what some have called "warranted assertability maneuvers" (e.g., DeRose 2009). Others have called them "pragmatic explanations of apparently semantic intuitions" (e.g., Davis 2007). For discussion of such views, see e.g., Unger (1975, 1984, 1986), DeRose (2009), Prades (2000), Rysiew (2001, 2007), Hawthorne (2004), Schaffer (2004b), Pritchard (2005), Brown (2006), Davis (2007), Pynn (2015), and Stoutenburg (Forthcoming). It is worth emphasizing that fictionalism is *not* a warranted assertability maneuver because it does not make any claims to the effect that we are warranted in asserting (potential) falsehoods because doing so allows us to convey truths. But, as noted below, it *is* a kind of pragmatic explanation of apparently semantic intuitions (albeit one that is distinctly different from many others of the same type).

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

There is, however, another general strategy available for resolving tensions in our use of certain expressions that has not yet been attempted in this domain: namely, *fictionalism*. Fictionalist views about a region of discourse can be provisionally characterized as accounts according to which claims made within that discourse are (typically) not best seen as aiming at literal truth, but rather should be interpreted as non-literal speech of some kind (cf. Eklund 2015).[8] And while such views might qualify as falling under (v) above, if it were construed very broadly so as to apply to *all* pragmatic explanations of the relevant variability in our use of knowledge-attributing sentences (and non-literal speech considered to be a pragmatic phenomenon),[9] they would be very special versions of it. For it is one thing to hold that pragmatic features of our use of knowledge-attributing sentences are more complex than one might have first thought, and quite another thing to hold that such sentences are generally used non-literally. And, of course, it is natural to think that if knowledge talk were largely non-literal, we would not *assert* knowledge-attributing (or -denying) propositions very frequently, if at all.

In this paper, I argue that such an alternative merits serious consideration. In Section 1, I survey a selection of fictionalist arguments before going on to explain, in Section 2, how they can be used to support the claim that knowledge-talk is largely non-literal. I also show how an epistemic fictionalist can avail herself of an additional argument concerning retraction behavior not generally employable by other types of fictionalist. In Section 3, I address several particularly pressing worries about the fictionalist's proposed approach that pertain to any account that would seek to posit an expansive divide between what sentences of the form "S knows p" are typically used to assert and their semantic contents. On the basis of this discussion, I conclude that the form of epistemic fictionalism outlined here deserves a spot at the table as a way of explaining the extensive variability in our use of knowledge-attributing sentences.

1. FICTIONALIST ARGUMENTS

Due to the wide variety of types of fictionalism, it is not possible to attempt anything like a comprehensive review of the styles of argument that have been proposed in favor of them

8 This characterization of what it means for a view to be fictionalist may strike some as insufficiently precise. However, given the variety of philosophical views that have been called "fictionalist," it is difficult to give both a more specific and yet still fully general account of what fictionalism is and what its distinctive commitments are. A more exact characterization differentiates two subtypes: *force fictionalism* and *content fictionalism*. Views of first type would have it that expressions within the relevant region of discourse are not generally used to assert (or otherwise illocute) any propositional content; rather, some other (perhaps perlocutionary) speech act is performed. Views of the second type would have it that expressions within the relevant region of discourse are generally used to assert (or otherwise illocute) propositional contents distinct from (and instead of) their semantic contents (cf. Eklund 2015).

9 For the purposes of this paper, I assume that it is. Because of this, I use "semantic content" and "literal content" more or less interchangeably, for ease of explication. It is worth noting, however, that some might take issue with this on the grounds that it may be possible for non-literal content to be semantic content, at least in certain cases (see Camp 2006a for a discussion of possible examples). As far as I can tell, however, accommodating this worry would affect only the letter, but not the spirit, of the view advanced here, as it could easily be restated so as to alleviate this concern. If it were, however, it might constitute a special version of approach (iii).

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

here. However, there are nonetheless a number of forms of argument that pertain to fictionalism as a general strategy rather than to any specific view. This section briefly surveys several of the most influential among them, with an eye to those that are likely to be of particular use to an epistemic fictionalist. These concern: (i) eliminativist arguments and error, (ii) thought experiments involving an Oracle, (iii) paradoxes, and (iv) analogies with uncontroversially non-literal discourse. Since the thesis defended here is intended to apply to our actual use of knowledge-attributing sentences, these arguments will be considered as they pertain to descriptive or "hermeneutic" forms of fictionalism only, rather than prescriptive or "revolutionary" forms.[10]

1.1 *Eliminativist arguments and error*

Perhaps the most familiar way of motivating descriptive forms of fictionalism about a region of discourse D proceeds by arguing that sentences within D are typically false, say, because the entities or properties to which they refer (for example, numbers, other abstract objects, possible words, and so on) don't exist, or because its characteristic predicates are uninstantiated or unsatisfied (nothing is right, nothing is good, and so on). It is then pointed out that this threatens to convict speakers who utter those sentences of massive error, which is generally regarded as highly implausible, especially if the linguistic practices in question seem to be successful. Fictionalism is then posited as providing an explanation for why these practices are effective independently of the (potential) falsity of the sentences in question (cf. Eklund 2015). After all, if the utterances in question aren't literal, then speakers needn't be thought to use the sentences in question to assert (or otherwise illocute)[11] e.g., the false propositions that they express; rather, they might even be thought to use them to assert truths.

1.2 *The Oracle*

This style of argument originates in a thought experiment proposed by John Burgess and Gideon Rosen in their seminal work, *A Subject with No Object*, and proceeds as follows:

> Finally, after years of waiting, it is your turn to put a question to the Oracle of Philosophy. So, you humbly approach and ask the question that has been consuming you for as long as you can remember: "Tell me, O Oracle, what there is. What sorts of things exist?" To this the Oracle responds: "What? You want the whole list? Look, I haven't got all day. But I will tell you this: everything there is is concrete; nothing there is is abstract. Now go away and don't bother me." (Burgess and Rosen 1997: 3)

The moral of this story is thought to become apparent when readers are asked to envision themselves as the protagonist of this fable and to reflect on the following two questions.

10 Please note that the structure of this section closely follows Eklund (2015).

11 I speak almost exclusively in terms of assertion in what follows, for simplicity; however, once again this restriction would have to be removed in a complete analysis.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

(1) Do you think that you, following your encounter with the Oracle, would persist in talking as before – that is, do you think that you would continue to use sentences that seem to commit you to the existence of abstract objects?
(2) If so, do you think that you would be using such sentences any differently than you were before? (Cf. Yablo 2000; Eklund 2015.)

To the extent that one expects that they would continue to use such sentences in the same way after meeting the Oracle (and that others would, too), one has gone some measure toward granting the fictionalist under discussion the ingredients that they need to cook up the following argument in favor of their view. For if, after returning from our chat with the Oracle, we would all go on using the relevant sentences as we were before, then just as we wouldn't be using sentences that seem to commit us to abstract objects entirely seriously – that is, literally – *after* speaking with the Oracle, we weren't using them literally *before* talking with them, either. In other words, insofar as we're inclined to think that we would most likely continue to use sentences that seem to commit us to abstract objects even with the (Oracle-given) information that there aren't any, and that nothing would seem all that different about our use of such sentences, then since we wouldn't be committing ourselves to positive claims about abstract objects post-Oracle, we weren't committing ourselves to positive claims about abstract objects pre-Oracle either – even if the sentences that we use are in fact so committal. Further, it is important to note that while this particular incarnation of the Oracle argument is designed to support descriptive fictionalist views about abstract object discourse in particular, it should be clear enough as to how it can also be used so as to argue for other kinds of fictionalism once suitably amended.

1.3 Arguments from paradox

These arguments proceed by drawing our attention to the fact that questions about whether certain sentences express truths can seem, depending on how one looks at them, easy or hard to answer (cf. Yablo 2000).
For example, the following can seem sound:

2 + 2 = 4
Therefore,
There are numbers

The problem, of course, is that this argument (and many others like it) can also seem unsound. After all, it is controversial – at least in some circles – as to whether there are numbers, and other abstract objects such as properties, possible worlds, moral facts, and the like. Because of this, in certain contexts, these are provocative claims, despite the fact that in most contexts, they are banal (cf. Dorr 2005).[12]

Here, too, the fictionalist can aid us in resolving the problem at hand. Typically, they'll say, when we use the sentences in question we use them non-literally so as to avoid

12 As Cian Dorr puts the point, sentences such as "2+2=4" and "There are numbers" are ordinarily used to assert boring, well-known truths (Dorr 2005: 2). But they aren't always; sometimes, they are used to assert exciting (at least to some), contentious claims.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

asserting their semantic contents and instead to assert something else altogether. However, in some situations (for example, those in which we are discussing whether there are such things as numbers, or other abstract objects such as properties, possible worlds, or moral facts), we do not speak non-literally when we use such sentences. Rather, we use them literally, and assert the propositions that they semantically express.

Thus, the fictionalist can give a neat, elegant, and reasonable-sounding solution to the puzzles presently under consideration. Questions concerning whether certain sentences express truths can seem easy to answer when we are speaking non-literally. However, questions concerning whether certain sentences express truths can seem hard to answer when we are speaking literally, such as when eliminativist arguments are in view. For such arguments seriously call into question whether the propositions semantically expressed by the relevant sentences are true, focusing our attention on their literal uses.

1.4 *Analogies with uncontroversially non-literal discourse*

A number of fictionalists, most prominently Stephen Yablo, have attempted to support their views by comparing our use of sentences in the target discourse with uncontroversial examples of non-literal speech and arguing that they share a number of core features. Yablo in particular has drafted up a fairly substantial list of such features, noting that while none of these similarities is anything like decisive, and that not all features mentioned are common to all examples, the cumulative effect is "nothing to sneeze at" (Yablo 2000: 21). Although there are far more examples on Yablo's list, the most important for our purposes concern silly questions, impatience, felt gaps, and unobtrusive non-literality (cf. Yablo 2000).

1.4.1 Silly questions and impatience

To borrow an observation from Zoltán Szabó, the best cases for descriptive forms of fictionalism about a given discourse always seem to rest on the feeling that certain philosophical objections somehow miss the point of that discourse – in other words, that they strike us as being silly questions (Szabó 2001: 378). For example, if a person is attempting to figure out, say, whether "2 + 2 = 4" is being used to express a truth, it seems rather beside the point (at least, in most contexts) to ask that person whether they believe that there are numbers. Somehow, it seems as if this question is not directly targeting what this sentence is typically being used to express. Indeed, the natural reaction to such a question (at least, in most contexts) is analogous to the natural reaction to questions like "If the pool is a mile long, could you run the length of it in under ten minutes?" "If the house is silent, why can I still hear the refrigerator?" and "How did the butterflies in your stomach get there? Did you eat them for lunch?" Such questions often cause people to become, as Yablo has pointed out, impatient, because they do not seem to be pertinent to what is under discussion. Or they strike people as silly, plausibly in the same kind of way that questions that treat non-literal utterances as literal do. This congruity can be appealed to in support of the view that the relevant sentences are, at least in ordinary discourse, used non-literally rather than literally. For, although the sorts of questions just mentioned are not usually considered to be relevant, or serious, there are contexts in which they are. How might this be explained? Fictionalists have a simple and straightforward answer to this question. These questions are relevant when the literal content of the relevant sentence

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

is in question. When what is in question is something else, however, it is not at all clear that they are, and it often seems as if they are not. Because the latter sort of context is the most common, these questions strike us as irrelevant, or silly, most of the time. But, because we are sometimes in the former sort of context, they do not always.

1.4.2 Felt gaps[13]

Non-literal speech may rely on a felt (albeit not usually explicitly processed) gap between literal and intended meaning, or expression and speaker meaning (cf. Camp 2006a, 2009). That is, it is plausible that, in order for an utterance to count as non-literal, language users must be in some sense (or on some level) aware of the distinction between the two, as well as of the connection between them, even if that awareness consists in having mere dispositions to think and act in certain ways, rather than, e.g., explicit or second-order beliefs about whether sentences are being used non-literally or not. For instance, we can expect speakers to feel compelled to articulate the difference between the literal and the intended meanings of their non-literal utterances if called on to do so, and to attempt a rough rational reconstruction of the connection between them (Camp 2006a, 2008, 2009).

So, what about sentences that fictionalists have targeted? Do speakers act in ways that suggest that they might be sensitive to the fact that such sentences often have both literal and intended meanings that differ from each other?

There is strong reason to think that the answer to this question is "yes," at least in some cases. Consider sentences that (appear to) make reference to possible worlds. As Yablo points out, such sentences can often be paraphrased with no felt loss of subject matter. For starters, "There is a possible world at which S" can be restated as "It is possible that S." And about the best that one can do with, "What if for every world, there's another more populous one?" is to hear it as the (not so bizarre) suggestion that you can always have one more thing (Yablo 2000: 21).

What's more, even if the sentences targeted by the fictionalist cannot be successfully paraphrased in such a manner, it is possible to argue for felt gaps in another way. After all, speakers are frequently very surprised to learn that particular sentences that they use plausibly carry with them certain commitments that they might at the same time find much harder to accept. This may well suggest a significant difference between what speakers are inclined to use certain sentences to convey and what the sentences do in fact express. To see the appeal of this way of thinking, try to remember the first time that you heard an eliminativist argument. Did its presuppositions regarding what the sentences in question express to some degree surprise you? But did you feel the intuitive pull of the argument nonetheless? If you are inclined to answer these questions with a "yes," and are inclined to think in addition that people are generally disposed to react this way to eliminativist arguments, it is reasonable to think that a discrepancy between speaker meaning and expression meaning might explain this.

13 This is a better way, in my view, of capturing what Yablo appears to be attempting to pick out with what he calls "paraphrasability" and "translucency" (Yablo 2000). It is perhaps worth noting, however, that Yablo, while using tests like paraphrasability, also sometimes stresses (in, e.g., Yablo 1998) that he does not know what it would be like to use sentences such as "There is a city of Chicago" more literally than he uses it now, whereas nothing about the arguments noted here relies on claims of this sort. (Thank you to an anonymous referee for registering this point.)

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

1.4.3 Unobtrusive non-literality

This argument works by attempting to remove an obvious objection along the following lines. Non-literal speech is not usually all that hard to detect. However, utterances of sentences that fictionalists target do not generally strike us as non-literal. Therefore, there is good reason to think that they are not, *pace* the fictionalist, typically used non-literally. One way to respond is to point out that non-literal speech is both pervasive and inconspicuous. Because of this, we typically do not notice it unless our attention is directed toward its non-literality.[14] And if this is true, we should not be surprised if some philosophically interesting discourse should turn out to be thoroughly permeated with non-literality. Examples are normally employed to substantiate this. Here is a particularly impressive list of potentially inconspicuous metaphors (or at least instances of non-literal speech) suggested by Yablo:

> [T]hey put a *lot of hurdles* in your path, there's *a lot* that could be said about that, there's *no precedent* for that, *something* tells me you're right, *there are some things* better left unsaid, *there is something* I forgot to tell you, viz. how to operate the lock, *nothing* gets my goat as much as chewing gum in class, *a lot* you can do for me, judge then of *my surprise*, let's roll out *the red carpet, the last thing I want* is to …, they have been rising in *my esteem*, I took her into *my confidence*, I'll take *my chances*, there's a *trace of sadness* in your eyes, a *growing number* of these leaks can be traced to Starr's office, she's got *a lot of smarts*, let's pull out *all the stops*, let's proceed along *the lines suggested above* … (Yablo 2000: 14)

Yablo's central point is that, if a such a variety of non-literal speech can be this unobtrusive, then maybe the (possibly even less conspicuous) non-literality of many of our other utterances can be too, perhaps depending to some extent on the kind of non-literality involved (one might think that some forms are more noticeable than others). That said, even the most flagrant forms of non-literality are arguably often not explicitly believed to be such in many cases (nor do people have the relevant second-order beliefs). For example, many cases of sarcasm are easily processed without one having to consciously attend to their non-literality (for example, certain utterances of "(So-and-so) is a fine friend") (cf. Camp 2006a).

It should be emphasized, however, that this argument – unlike others just canvassed – is probably best interpreted as providing indirect rather than direct support for fictionalism, on the grounds that it works by aiming to remove a very natural worry that one might have about the approach. This is because it, considered alone, does not propose to offer evidence for the claim that in using certain sentences we generally aim to convey something other than those sentences' literal (i.e., semantic) contents, unless one thinks that unobtrusiveness is among the defining features of non-literal speech, which seems highly implausible on the grounds that literal speech is also unobtrusive (it would have to be, after all, if non-literal speech is). However, it may help to make arguments that do attempt to render such evidence much more plausible and, to the extent that it does, may prove an indispensable tool for fictionalists just the same.

14 While one might wonder whether "felt gaps" and "unobtrusiveness" might be in tension with each other, it seems that much non-literal speech is Janus-faced in this way.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

2. APPLYING FICTIONALIST ARGUMENTS IN EPISTEMOLOGY

2.1 *Epistemic fictionalism: a formulation*

In order to show how the arguments just canvassed can be used to defend a fictionalist view in epistemology, in what follows I supply a sketch of what such an account might look like, the resulting position being a form of *content epistemic fictionalism* (henceforth simply "epistemic fictionalism" for the purposes of this paper). Such an approach would have it that expressions within the relevant region of discourse – namely, knowledge-discourse – are generally used to assert (or otherwise illocute) some propositional content distinct from (and instead of) their semantic contents.[15] I then explain how the fictionalist arguments discussed above can be used to support it.

As noted in the introduction, philosophers have long sought to explain an array of kinds of variability in our use of knowledge-attributing sentences. This task, however, has proved notoriously difficult. The fictionalist's main insight is that much of this variability may be due to the fact that we often use knowledge-attributing sentences non-literally, to assert (or otherwise illocute) a variety of propositions. To see how this might help, imagine that we grant that skeptical arguments are possibly sound; that perhaps they succeed in establishing that no one knows much, if anything, at all. From this it would follow, given a standard semantics, that sentences of the form "S knows p" (and the like) usually, if not always, express false propositions. Does this mean that we are not warranted in our use of such sentences? No, says the fictionalist. For, we do not normally use a sentence of the form "S knows p" to assert the proposition that it semantically expresses. Rather, we use it non-literally in such a way so as to avoid asserting its semantic content and to assert some other content (consistent with the proposition that S does not know p) altogether.[16]

What might this alternative content be? Since non-literal speech is often deeply context-sensitive, we should expect the answer to this question to be more or less context-sensitive as well, depending on the pragmatic machinery involved (given that some pragmatic phenomena tend to be more deeply context-sensitive than others). Because of this, it may be very difficult (if possible at all) to completely spell out this potential alternative content in full generality, much less in any specific case (especially if one takes seriously the possibility that the propositional content that a non-literal utterance of a sentence of the form "S knows p" is used to assert is often indeterminate across a range of (determinate) propositions, as is frequently the case with non-literal assertions).[17] Typically asserted

15 For more on the distinction between force and content fictionalism, please see fn. 8.

16 Note that if this is right, it does not matter whether skeptical arguments are sound – at least not if they target the semantic (i.e., literal) content of knowledge-attributing sentences, viz., knowledge claims, rather than the content that they are used to pragmatically convey. Because of this, a fictionalist need not commit herself to the claim that skeptical arguments are sound, and that knowledge-attributing sentences are false. The fact that both are serious possibilities is sufficient.

17 For example, if the relevant pragmatic machinery is that of metaphor, as it is in e.g., Yablo (2001) (a possibility that I pursue in Chung (Forthcoming)), then there is strong reason to suspect that the content may be indeterminate (cf. Walton 1993; Soames 2002; Camp 2006b). For a more detailed explanation of why this is reasonable, please see section 3.3 below. However, other types of pragmatic machinery may yield far more determinate content – say, if the relevant machinery pertains only to relevant or salient alternatives (along the lines of, e.g., Rysiew 2001). For the purposes of this paper, I propose to remain neutral between these disparate possibilities to the extent that I can, as

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

propositions may include those along the lines of: that S truly believes p, that S is justified in believing p, that S believes p as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to p, that S can assert p, that S can rely on p in practical reasoning, and so on – in other words, propositions that many epistemologists have argued are associated with knowledge-attributing sentences.[18] However, it should be emphasized that these may not even come close to exhausting the propositions that we use sentences of the form "S knows p" to assert. Indeed, the fact that there could be an expansive and elusive array of such propositions might at least partially explain many other kinds of variability in our use of knowledge-attributing sentences (including, as mentioned above, the kinds of variability that occur as a result of changing what is at stake or the salience of alternative possibilities, as well as others in addition).[19]

This account has two potentially attractive features worth singling out at this time, as they are akin to just those features that are commonly thought to make fictionalist views appealing in general. First, one can endorse it in conjunction with classical invariantism regarding the semantic content of "knows." This carries with it two distinct benefits: one can maintain (i) a simple semantics for "knows" (thus avoiding many of the problems faced by semantic explanations of variability in our use of knowledge-attributing sentences, such as contextualism) and (ii) a simple metaphysics of knowledge (thus avoiding many of the problems faced by other forms of invariantism, such as subject-sensitive invariantism). Second, because it interprets much of our ordinary knowledge-talk in terms of the more widely recognized phenomenon of non-literal speech, it need not require the introduction of any new pragmatic machinery or special rules (thus avoiding many of the problems faced by a number of other pragmatic explanations of variability in our use of knowledge-attributing sentences).[20]

Of course, as in the case of other types of fictionalism, whatever its problem-dissolving potential, one might think that this view has no chance. After all, knowledge-talk – just like talk about numbers, other abstract objects, possible worlds, morality, and the like – does not generally strike us as non-literal in any way. However, despite this, the form of fictionalism outlined here can be supported by the same sorts of arguments that are usually employed in the service of fictionalist views. In addition, it can be supported by another argument, not universally employable by other types of fictionalists, concerning retraction behavior. I thus take there to be excellent reason to think that an epistemic fictionalism along the lines I sketch above is likely to be at least as plausible

my aim here is to motivate content epistemic fictionalism as a general approach, rather than any particular version of it. I would like to thank an anonymous referee for encouraging me to clarify this.

18 The candidates provided here are mentioned because, for each, there is a wealth of literature arguing that it, e.g., is a truth-condition for sentences of the form "S knows p," and among the fictionalist's key insights is that many of these supposed, e.g., truth-conditions may rather be part of the multitude of propositions that we in fact non-literally assert when we use sentences of the form "S knows p."

19 I discuss this in more detail in Chung (Forthcoming).

20 Whether it will in fact require the introduction of any new pragmatic machinery will depend on the specifics of the view. If the relevant pragmatic machinery is metaphor, hyperbole, or loose speech, for example, no new pragmatic machinery will be needed (cf. Schaffer 2004b). If it is something else, however, this may not be so.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

as other forms of fictionalism, if not more so.[21] The plan, then, is to devote the remainder of this section to illustrating how all this can be done.[22]

2.2 *Skeptical arguments and error*

No matter what, skeptical arguments – like eliminativist arguments – threaten to convict speakers of having made some sort of mistake. For although it seems as if sentences of the form "S knows p" are used to express truths, a number of initially appealing skeptical arguments suggest otherwise. It looks, then, as if we must be wrong about something: either sentences of the form "S knows p" are usually, if not always, used to express falsehoods, or skeptical arguments, despite appearances to the contrary, are (at least in many cases) unsound. Epistemic fictionalism offers a novel and tenable explanation for why we want to insist on saying things like "S knows p" even if we can't refute the skeptic, and for why these practices are effective independently of the (potential) falsity of the sentences in question. Further, we can see that there is a comparison between knowledge-talk and non-literal speech worth drawing here, as we often credit speakers with engaging in non-literal speech when doing otherwise would be to charge them with possibly falling into error by making a false assertion, particularly if we think that they are (a) also in some sense or on some level aware of this and (b) clearly attempting to state a truth.

That said, epistemic fictionalism, like all other forms of fictionalism, does appear to require the controversial psychological claim that speakers are at the same time generally either in some sense or on some level unaware of (or perhaps even at times flat-out wrong about) the fact that knowledge-attributing sentences are typically used non-literally. This is where the argument from unobtrusive non-literality can come to the rescue. For this sort of thing is far from unusual. Indeed, it is plausible to think that non-literal speech normally goes (in some important way) unnoticed by us, as much of our ordinary speech exhibits this quality. To see this, one need only reflect on the variety of kinds of non-literal speech (including, but not limited to, the kinds of non-literal speech mentioned earlier), and on how frequently we engage in them. What's more, not only is non-literal speech regularly not explicitly believed to be such (nor is it the case that people regularly have second-order beliefs to the effect that they are or are not speaking non-literally), it is arguably often misidentified as literal in addition (especially when the locutions in question are sufficiently common). If this were not so, it would not be possible to debate whether a given utterance was literal or non-literal. Thus, it should not be all that surprising if it turns out that much of our knowledge-talk is non-literal as well, and that people largely lack explicit (or second-order) beliefs to the effect that it is, or even frequently have mistaken explicit (or second-order) beliefs to the effect that it isn't. So, that speakers are

21 For ease of exposition, particular arguments employ examples of metaphor, as I consider that option to be at least as attractive as the rest, and because it is easy enough to see how they could be recast so as to employ other kinds of examples of non-literal speech.

22 It is worth noting that none of this is to say that *only* epistemic fictionalism can account for the phenomena discussed below. Rather, my goal is more modest: to show that epistemic fictionalism can be supported by the same arguments that generally motivate fictionalist views. The question of whether it offers a better explanation of the relevant phenomena than do the existing approaches canvassed in the introduction is thus left open for present purposes, though a full defense of the view would require it.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

largely in some sense (or on some level) not cognizant of, and even at times incorrect about, the fact (if it is one) that sentences of the form "S knows p" are typically used non-literally needn't be thought to pose all that serious, or special, a problem for the fictionalist, if it were thought to pose a problem at all.

2.3 *The Oracle part II*

Here we are, back in line, waiting to talk to the Oracle. Finally, after years more of waiting, it is our turn to put a question to them once again. Fantastic! We humbly approach and ask the *other* question that has been consuming us for as long as we can remember: "O Oracle, could you please answer just one more question for us? What, if anything, is wrong with skeptical arguments?" To this the Oracle responds: "You again? I don't know why you waste my time with these questions. There is absolutely nothing wrong with skeptical arguments. Now go away and don't come back!"

Dejected, we withdraw. The Oracle isn't pleased with us, but at least we received an answer to our question. No matter what, it seems as if the Oracle has given us: (a) knowledge that skeptical arguments are sound, (b) sufficient reason to believe that skeptical arguments are sound, or (c) sufficient reason to be agnostic about whether skeptical arguments are sound. From this, we can infer that we either have (a) knowledge that many sentences of the form "S knows p" express something false, (b) sufficient reason to believe that many sentences of the form, "S knows p," express something false, or (c) sufficient reason to be agnostic as to whether many sentences of the form, "S knows p," express something false. Of course, one might think that skeptical arguments, if sound, show that we have no knowledge whatsoever, and in that case the (a) option above is out (indeed, it is arguably out even if we are agnostic about whether skeptical arguments succeed). That's okay, though – we don't really need it. In fact, we don't really need the (b) option either. So long as we have the (c) option, we have all that the epistemic fictionalist requires.

Let's imagine that all of this happened. We are now in the position to reflect on the following two questions once again.

(1) Do you think that you, following your encounter with the Oracle, would persist in talking as before – that is, do you think that you would continue to use sentences of the form "S knows p"?
(2) If so, do you think that you would be using such sentences any differently than you were before?[23]

23 As a note, even if one is inclined to think that we might feel a little different about our uses of such sentences, that would not undermine the Oracle argument. For, we might feel different in some way – say, because we might become more sensitive to the false literal interpretations of the relevant sentences – but not about whether we were committing ourselves to their contents in the first place. Compare a case of metaphor: If I were to become consciously aware of the fact that a given phrase is used metaphorically (say, "grind out" as in, "I'm trying to grind out this paper"), I might feel a little different when I use it in the future because this fact would have been made salient to me and I would think about it, but surely I wouldn't feel different about whether I was committing myself to the literal content of the relevant expression.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

If you are inclined to answer these questions with a "yes" and a "no" respectively, then the epistemic fictionalist just might be able to convince you that their view is worthy of your consideration. For if we would continue to use sentences of the form "S knows p" despite the fact that we would either (b) believe that they express falsehoods, or (c) be agnostic as to whether they express falsehoods, and nothing would seem all that different about our use of such sentences, then since we wouldn't be using the relevant sentences to commit ourselves in the way under discussion post-Oracle (after all, we would either (i) believe that many sentences of the form, "S knows p", express false propositions post-Oracle, or (ii) be agnostic as to whether sentences of the form, "S knows p," express false propositions post-Oracle), then we weren't using them to so commit ourselves pre-Oracle, either.

2.4 *A paradox about knowledge*

This problem begins by drawing our attention to the fact that questions about what sorts of things are known can seem, depending on how one looks at them, easy or hard to answer.

For instance, the following argument can seem sound:

> S knows p (i.e., where p is a proposition about the external world, as in: S knows that here is a hand.)
> Therefore,
> There is a proposition about the external world that S knows

But so can skeptical arguments that purport to show that such things are not (and indeed cannot be) known. And, once one starts thinking about those, the above argument can start to seem unsound.

Here, too, the fictionalist can aid us in resolving the problem at hand. Typically, they'll say, when we use a sentence of the form "S knows p", we use it non-literally in such a way as to avoid asserting its semantic content (the proposition that S knows p), and instead to assert, for example, that S can rule out all relevant alternatives to p (possibly among other things). However, in some situations (for example, those in which we are discussing whether skeptical arguments are sound, or in which stakes are high or skeptical arguments or alternatives have somehow been made salient), we do not speak non-literally when we use such sentences. Rather, we use them literally, and assert the propositions that they (semantically) express.

Thus, the epistemic fictionalist can also give a neat, elegant, and reasonable-sounding solution to the puzzle under consideration. Questions concerning what we know can seem easy to answer when we are speaking non-literally. However, questions concerning what we know can seem hard to answer when we are speaking literally, such as when skeptical arguments are in view. For, such arguments seriously call into question whether the propositions semantically expressed by knowledge-attributing sentences are true, focusing our attention on their literal uses.[24]

24 One might wonder why the question remains hard to answer when we speak literally. After all, on the fictionalist's view, shouldn't it be obvious to all competent speakers that knowledge-attributing sentences are straightforwardly false? To see why not, consider what a fictionalist about mathematical

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

2.5 *More silly questions and impatient reactions*

Recall our silly questions from earlier. Can an epistemic fictionalist argue that certain queries involving knowledge also strike us as somehow beside the point, or provoke impatient reactions? Once again, the answer is yes. For if a person is contemplating whether a sentence of the form "S knows p" is being used to express a truth, it seems similarly beside the point (at least, in most contexts) to ask that person whether, for example, they believe that people can rule out the possibility that they are brains in vats (even if p is a proposition about the external world), or whether they think that people can ever substantiate their claims in a non-viciously-circular fashion. Such questions often cause people to become, as Yablo might say, impatient, because they do not seem to be relevant to what is under discussion. Or, they strike people as silly, plausibly in the same kind of way that questions that treat non-literal utterances as literal do.

However, even though the sorts of questions (that is, those related to skeptical arguments) just mentioned are not usually considered to be relevant, or serious, there are contexts in which they are. How might this be explained? The proponent of epistemic fictionalism has a simple and straightforward answer to this question. These questions are relevant when the literal (semantic) content of a sentence of the form "S knows p" is what is at issue – that is, when what is in question is whether S knows p. However, when what is at issue is something else (for example, whether S can rule out all relevant alternatives to p), it is not at all clear that they are, and it quite often seems as if they are not. Because the latter sort of context is the most common, these questions strike us as irrelevant, or silly, most of the time. But because we are sometimes in the former sort of context, they do not always.

2.6 *Felt gaps*

Speakers also act in ways that suggest a possible perceived gap between what sentences of the form "S knows p" express and what speakers use them to express. To see this, we might observe that although it is admittedly difficult for speakers to spell out what they used a sentence of the form "S knows p" to pragmatically convey (as opposed to what it semantically expresses) it is not completely beyond their ken. Rather, the English language arguably has some devices available for this purpose. For instance, one can clarify one's knowledge attributions by saying things like: "They know, in effect," "They know, with reasonable assurance," "They know full well," "They pretty much know," "They know well enough," "They know, for certain," "They know, without a doubt," "They know, assuming that q," and "They don't literally know"

---

discourse might say about "2+2=4" and "There are numbers": the reason that it is not obvious to competent speakers that these sentences are false is that the question of whether there are abstract objects such as numbers is a difficult one. However, we don't use such sentences to commit us to their semantic contents – and because of this, it does not matter whether they are true or false. Likewise, an epistemic fictionalist could maintain that the reason it is not obvious to competent speakers that knowledge-attributing sentences are false is that the question of whether or not we know anything is a difficult one – but that it similarly does not matter whether they are true or false. (In connection with this, compare fn. 16 above.) Many thanks to an anonymous referee for raising this question.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

(cf. Hawthorne 2004: 105).[25] Further, even if the sentences targeted by the fictionalist cannot be successfully paraphrased in such a manner, once again it is possible to argue for felt gaps in another way. After all, speakers are often very surprised to learn that sentences of the form "S knows p" plausibly carry with them certain commitments that they may at the same time find much harder to accept. And this may well suggest a significant difference between what speakers are inclined to use certain sentences to convey and what the sentences do in fact express. To see the appeal of this way of thinking, try to remember the first time that you heard a skeptical argument. Did its presuppositions regarding what the sentences in question express to some degree surprise you? But did you feel the intuitive pull of the argument nonetheless? If you are inclined to answer these questions with a "yes," and people are generally disposed to react this way to skeptical arguments, it is reasonable to think that a discrepancy between speaker meaning and expression meaning might explain this.

2.7 *Retraction behavior*

Another way in which knowledge-talk can be argued to bear similarity to uncontroversially non-literal speech concerns retraction behavior. It is commonly acknowledged that skeptical arguments tend to put pressure on people to "take back" or "retract" their uses of knowledge-attributing sentences. Upon being presented with such arguments, many are inclined to assent to the possibility that they spoke falsely when they previously uttered sentences of the form "S knows p." Many also find themselves wanting to say things like "S didn't really know p" with certain kinds of emphases. Despite this, however, many seem positively disinclined to change their knowledge-attributing practices much, if at all.

These patterns are just what one should expect if epistemic fictionalism were correct. For it seems as if hearers are always – or, at least, almost always – entitled to insist on a sentence's literal interpretation, even if doing so would be extremely uncooperative (cf. Camp 2006a: 296–7). And, in the face of such challenges, although speakers are expected to withdraw their non-literal utterances, they are not typically expected to feel as if they shouldn't have spoken non-literally, or that they should refrain from speaking non-literally in the future. This is even true of fairly clear cases of non-literal speech.[26] Consider the following:

A: I have butterflies in my stomach.
B: That's absurd. All you've had to eat all day was oatmeal, and they surely could not have gotten there any other way!

25 While some might find it problematic in this context that "knows" is used in these purported clarifications (and thus, that they are not sufficiently informative), the force of this worry is blunted if one is inclined to think that what they possibly indicate is a hesitance to commit to the literal contents of sentences of the form "S knows p" (even if they do not specify any alternative contents all that precisely). Thank you to an anonymous referee for encouraging me to address this concern.

26 I say "fairly clear" because some might argue that this example is idiomatic. Since I have chosen it for the sake of brevity, I encourage those who are concerned about such matters to substitute a novel piece of non-literal speech of their choice.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

In this sort of exchange, A can't typically simply dismiss B's literalist challenge by saying, "You're wrong. I do have butterflies in my stomach," or even, "You're wrong. I am really nervous" (cf. Camp 2006a: 296). Rather, there seems to be some pressure on A to retract their utterance by acknowledging that what they said (that they have butterflies in their stomach) was false, even if what they asserted (roughly, that they are nervous) was true.[27] What's more, A can acquiesce to this pressure without having the strong sense that they shouldn't have said what they did, and without feeling as if they should stop using this sentence (and others like it) in this way. This is not what we would expect if A believed themselves to have made a false assertion. Thus, it seems reasonable to say that what A is retracting in this case is the semantic content of their utterance, or what they said, rather than their claim, or what they asserted.

We can expect something similar to happen in this conversation:

A: I know p.
B: No, you don't. After all, (insert your favorite skeptical argument).

Here, too, there seems to be some pressure on A to retract what they said by acknowledging that it at least may have been false, even if what they asserted was true. Also, A can very often acquiesce to this pressure without having the strong sense that they shouldn't have said what they did, and without feeling as if they should stop using such sentences in this manner.

Interestingly, it is worth pointing out that the fictionalist can also account for what is going on when speakers "stick with" their utterances (even after having been presented with skeptical arguments), which, as many commentators have pointed out, tends to happen from time to time (cf. Cohen 2000b; Hawthorne 2000; Davis 2007, 2015). What they are sticking with is what they asserted (what they used a sentence of the form "S knows p" to convey), not what they said (what the sentence that they used expresses). However, although epistemic fictionalism allows for this, it predicts that it should be harder to pull off, which seems the right result. For, although speakers who engage in non-literal speech typically face significant pressure to retract their utterances in the face of literalist challenges, this is not always the case. Let us again consider a fairly clear case of non-literal speech. If someone says:

27 Now, they might hold that what they meant, or asserted (presumably, something like the proposition expressed by "I am nervous") was true, and they might make a point of stating this intended meaning literally and explicitly. Some might complain that there is some dissimilarity, then, between this sort of case and cases like "S knows p" (for in the latter case, speakers cannot be much more explicit than they already are). Although this is acknowledged (with the caveat that there are some ways that speakers try to make themselves more explicit when faced with such challenges, as discussed above (see also Hawthorne, 2004)) it is worth pointing out that there is the distinct possibility that a language may lack the expressive resources necessary for speakers to spell out what they meant by their non-literal utterances in every case (cf. Walton 1993; Camp 2006b). Thus, the fact that there may be no possible paraphrase of what a speaker meant by a non-literal utterance does not show that the speaker must have asserted what they said. Moreover, this is also what one should anticipate if speakers are largely in some sense or on some level unaware of the fact that they are typically speaking non-literally when they use sentences like "S knows p."

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

"Bill's a bulldozer,"

although a literalist challenge along the lines of,

"No, he's not; he's a man"

always seems in some sense legitimate, it can nonetheless be dismissed, rather than engaged with, if other conversational norms mandate it. For example, if one were in a hurry to finish making their point before Bill had a chance to overhear, they might respond to such a challenge by saying:

"C'mon; sure he is. After all, can you imagine anyone pushier than he?"

This form of argument is not generally employable by fictionalists. I expect that fictionalists about mathematical discourse, for example, would have a harder time convincing people – especially ordinary speakers – to consider retracting their uses of "2 + 2 = 4" or even "There are numbers." I anticipate that this would also be true of ordinary speakers' uses of sentences that appear to make reference to other abstract objects such as properties, as well as moral sentences.[28] Modal fictionalists might find such an argument useful, however – at least to the extent that it is plausible to think that those who go in for possible-worlds talk would be happy to assent to the claim that there aren't really possible worlds. For now, I will consider it an open question as to whether others could as well. For the principal point of the foregoing discussion is that retraction behavior can be used to motivate fictionalist views (insofar as the relevant patterns are similar to those that we observe with respect to uncontroversial or less-controversial examples of non-literal speech). And, it promises to offer a significant measure of support for epistemic fictionalism in particular.

3. GENERAL WORRIES

3.1 Assertion and semantic content

In *Knowledge and Lotteries*, John Hawthorne registers a variety of concerns about pragmatic explanations of variability in our use of knowledge-attributing sentences, the most relevant of which – as far as the fictionalist's proposal is concerned – will be addressed here. The first pertains to what Hawthorne describes as being "subtle issues about the nature of assertion." Many defenders of pragmatic explanations of variability in our use of knowledge-attributing sentences maintain that although such sentences very often, if not (nearly) always, express falsehoods, we are frequently warranted in asserting them because doing so allows us to convey truths. This, however, prompts the question of why (given that we are in some sense or on some level aware that knowledge-attributing sentences may express falsehoods) we should be thought to assert their semantic contents at all.

28 As a note, this might just be due to the fact that there are no non-philosophical challenges to, e.g., arithmetical claims, whereas there are perfectly ordinary challenges to knowledge claims, which may make us more tolerant toward the knowledge-skeptic. Thank you to Zoltán Szabó for this insight.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

I am inclined to agree that such considerations put pressure on pragmatic theorists – including fictionalists – to endorse a position on which the relevant semantic contents are not asserted. However, as Hawthorne points out, this in turn raises two pressing questions. First, if "knows" is typically used non-literally, then why hasn't our semantics evolved to reflect this? After all, if the customary use of a sentence s is to assert a certain propositional content, (almost) invariably, that content becomes the literal (semantic) content of s (cf. Szabó 2001: 305; Hawthorne 2004: 123). Second, if "knows" is generally used non-literally, from where could its alleged semantic content derive? Any pragmatic theorist who goes this route needs to tell a story about where knowledge-attributing sentences get their literal meanings, if not from facts about what they are customarily used to assert (Hawthorne 2004: 122–3). Although Hawthorne doubts that satisfactory responses to these questions can be supplied, I am more optimistic, and will provide answers to each.[29]

Regarding the first question, two principal considerations are relevant. The first concerns the ubiquity of skeptical arguments and the second, the frequently indeterminate nature of non-literally conveyed propositional content.

3.2 *The ubiquity of skeptical arguments*

As various commentators have pointed out, one doesn't need (seemingly) far-fetched scenarios involving extended dreams or brains-in-vats to cast doubt on a vast number of our everyday claims to knowledge. Much of the time, far more mundane suggestions will do.

To illustrate this, here are a few examples, originally due to Jonathan Vogel:[30]

A: Do you know where your car is parked?
B: Yes, it's parked on Elm Street.

A: Do you know that it hasn't been stolen since then?

A: Do you know who the current President of the United States is?
B: Yes, it's Barack Obama.

A: Do you know that he hasn't had a fatal heart attack in the last five minutes?

A: Do you know where I can get a good hamburger?
B: Yes, there's a luncheonette several blocks from here.

A: Do you know that a fire hasn't just broken out there?

In each of these cases, it seems as if A is calling into question B's claim to knowledge without invoking any sort of bizarre skeptical hypothesis. Further, such challenges are easy to

29 It is perhaps worth mentioning that although Hawthorne's objections come from a discussion of epistemology, they concern fictionalism generally. (Thank you to an anonymous referee for noting this.) As a result, several of the considerations mentioned below in response may also generalize beyond the domain of epistemology as well.

30 Vogel 1990b; they have been appropriately updated.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

construct. The formula is simple: all one has to do, when presented with a sentence that (at least appears to) attribute knowledge, is figure out one proposition that it entails that we would be intuitively disinclined to take ourselves to know.[31] Once this is done, one has the weapons needed to mount a skeptical attack. What's more, ordinary speakers can, and do, attack each other's (apparent) knowledge attributions in this sort of way. And though such challenges often strike people as being antagonistic, they just as often do not (at least in any straightforward fashion) seem to be completely changing the subject.[32] This would be extremely puzzling if sentences of the form "S knows p" expressed propositions with less demanding truth conditions, as such challenges would seem to miss the point almost entirely if this were so.

Another upshot of these considerations is that we often seem to act as if sentences of the form "S knows p" have more demanding truth-conditions than our use of such sentences in fully (or at least highly) cooperative (i.e., non-antagonistic) contexts might suggest. This gives us reason to suspect that the semantic contents of these sentences cannot be read in any straightforward way off facts about our standard uses of them alone, and especially not off facts about our uses of them in non-antagonistic contexts to the exclusion of all others.[33] Rather, to paraphrase Elisabeth Camp, antagonistic contexts are not only real but also sufficiently pervasive and illuminating so as to warrant inclusion in our theoretical deliberations instead of being set aside as abnormalities (or, as I will add, distinctively philosophical oddities). Indeed, as Camp points out, it may well be the case that our linguistic behavior in antagonistic contexts is more pertinent to the task of determining semantic content than is our linguistic behavior in non-antagonistic contexts, as conversations in antagonistic contexts may reveal conversational rules (as opposed to mere conversational conventions) which can be occluded by charity in highly cooperative contexts. If so, then the raising of standards stands to be more suggestive of semantic content than the lowering of standards. As Camp writes:

> When a speaker departs from or goes beyond conventionally encoded meaning … she speaks, as Grice says, "'under license' from other participants" (1989, 45). 'Accommodation' is the willingness (and ability) to grant such a license, by playing along to the best of one's abilities in determining the content of the speech act being undertaken. It is easy to ignore the optional nature of accommodation when attending to cooperative contexts, where there is so much pressure to be maximally charitable and to facilitate the conversational flow. Indeed, even in antagonistic contexts it can take considerable cognitive effort to notice the respects in which an utterance is open to multiple interpretations, and to suppress the speaker's manifestly intended meaning. In keeping with this near-automaticity of accommodation, Lewis frames the rule for accommodation for presupposition in terms of the default case of cooperation:
>
> If at time t something is said that requires presupposition P to be acceptable, and if P is not presupposed just before t, then – *ceteris paribus* and within certain limits – presupposition P comes into existence at t.

31 Or nearly enough entails (cf. Hawthorne 2004: 5, fn. 13).

32 Although they admittedly do on some occasions, at least in some way and to some degree. A complete analysis would need to fully account for these complexities; however, I will not attempt to do so here.

33 This terminology is due to Elisabeth Camp. For more details, see Camp (2013).

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

> Interpretively antagonistic conversational contexts fall precisely within the terrain of the *ceteris paribus* clause, however. In such contexts, hearers are prepared to pedantically insist on the most literal and strict interpretation of the speaker's words, forcing her to articulate her meaning explicitly and precisely if she wants it to be registered on the conversational scoreboard at all. But such contexts are still real, and they are pervasive enough and illuminating enough that they deserve to be included in our theoretical considerations and not just dismissed as aberrations. Lewis observes that "for some reason raising of standards goes more smoothly than lowering." The reason, I suggest, is that the raising of standards at least typically hews more closely to the standards encoded by the semantic content itself. (Camp 2013: 25–6)

Thus, in Camp's view, antagonistic contexts are theoretically revealing because "the fact that one or both parties is prepared to be only minimally cooperative, working to conversational rule, helps to reveal what the conversational rules *are*, in a way that is obscured when the conversational operation is maximally lubricated with convivial charitability. To shift metaphors, draining away the waters of charitability reveals underlying structure where we might otherwise see only an undifferentiated pond" (Camp 2013: 29).

Granted, such speculations require more argument than I have provided here if they are to present as fully compelling. However, the fact that they are at minimum serious possibilities should suffice to support epistemic fictionalism as a live option. If it could be shown that our uses of expressions in antagonistic contexts really were more pertinent as regards determining semantic content, then this would do more than just substantiate the claim that epistemic fictionalism is worth pursuing further – it would actually give epistemic fictionalism an edge over its, e.g., semantic competition.[34]

3.3 Indeterminacy

As noted above, the propositional content that a non-literal utterance is used to assert is – in certain cases – frequently indeterminate across a range of (determinate) propositions, and because of this, it is plausible to think that, given epistemic fictionalism, a sentence of the form "S knows p" may not generally be used to assert something (more or less) determinate. Rather, the content that a sentence of the form "S knows p" is typically used to assert may usually be to some degree indeterminate over the set of propositions: {that S truly believes p, that S is justified in believing p, that S believes p as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to p, that S can assert p, that S can rely on p in practical reasoning, …}. Thus, although it is held that we have some idea as to what such sentences are being used to assert, according to the epistemic fictionalist, it may be the case that we can't quite get a handle on precisely what it is.

Why should we agree, along with the fictionalist, that the relevant content may be frequently indeterminate in this way, however? Three considerations are worth noting in connection with this question. First, speakers' intentions are often insufficiently specific to determine, with any precision, just what proposition or propositions they are asserting, especially when they are speaking non-literally (cf. Walton 1993; Soames 2002: 134–5, *inter alia*). Just what does a speaker intend to assert when they say, for instance, "Bill

34 Many thanks again to an anonymous referee for encouraging me to clarify this.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

is a bulldozer"? Perhaps they intend to assert that Bill is inconsiderately forceful, that Bill is insensitive to obstacles, both, or neither (as they may intend to assert something else about Bill altogether). Likewise (and for example) it may be indeterminate as to whether a speaker intends to assert, in saying "S knows p," the proposition that S can rule out all relevant alternatives to p, or almost all, or just the ones the speaker cares about, or perhaps something completely different. Further, even if assertion is, e.g., a social phenomenon, similar questions will arise pertaining to what a speaker is committed to when they use a sentence non-literally. Second, the burden of proof is on those who would maintain that some asserted content is determinate to specify what that content is, or at least give some reason for thinking that it must be so. For if one can't spell out just what a sentence is being used to assert, why should one think that it is being used to assert a determinate propositional content at all? If this challenge is not met, it seems as if we are within our rights to assume that the content that a sentence is being used to assert is not determinate.[35] Third, and finally, if the contents that sentences of the form "S knows p" are used to assert were generally indeterminate, then this would explain why epistemologists have struggled to the degree that they have in their quest to specify the truth or assertability conditions of sentences of the form "S knows p," and the fact that while many theories have proved remarkably plausible with respect to certain cases, they have proved exceedingly implausible with respect to others, for indeterminate things are, by their nature, elusive.

3.4 Blocking semantic modification

We are now in a position to give an answer to the first question posed at the beginning of this section. For it should be clear that (at least) two things might have prevented non-literal meanings from becoming literal meanings in this instance: (i) non-literal assertions of sentences of the form "S knows p" are often used to assert indeterminate propositional contents (thus making it the case that there is nothing sufficiently determinate to semantically encode) and (ii) our ability to effectively wield commonplace skeptical arguments (such as those discussed under the heading "the ubiquity of skeptical arguments" above) depends on the literal meaning of "knows" being such that sentences of the form "S knows p" have highly demanding truth conditions – perhaps so highly demanding that they are (nearly) impossible to meet.

Though I take the first of these points to require no additional argument, one might understandably wonder why the second should be endorsed. The short answer is that skeptical arguments – including those that we employ on a regular basis (and not just in the philosophy classroom) – can succeed only if a case can be made for the claim that the subject in question does not know a given proposition; however, if knowledge-

35 One worry about this argument proceeds along these lines: surely we can use an expression to say something definite, even if we cannot say (in informative, independent terms) what this something is. The fictionalist can grant this, however, but note that part of what is going on here concerns the fact that if it is already acknowledged – given fictionalist arguments – that what is often asserted when a speaker utters a sentence of the form "S knows p" with declarative force is possibly something other than that S knows p, then in the absence of arguments that aim to establish that it is indeed something else in particular, there is no reason to think that a determinate propositional content is asserted. Many thanks to an anonymous referee for registering this worry.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

attributing sentences did not have demanding truth conditions, such cases would be difficult, if not impossible, to make. For skeptical arguments can perhaps only have real bite when the literal meaning of "knows" is such that it is very difficult for sentences of the form "S knows p" to be true. Otherwise, challengers would be open to allegations of having made non-sequitur remarks, and although (as noted earlier) speakers can sometimes "stick with" their non-literal assertions even in the face of skeptical arguments, it does not seem as if they can accuse those who employ them of completely changing the subject. And while other theories on the market may be able to account for these phenomena as well, it is clear that epistemic fictionalism too offers a plausible way to explain them.[36]

3.5 Knowledge-talk and semantic content

This brings us to the second of Hawthorne's questions. Explaining exactly how the semantic content of an expression gets to be its semantic content is no easy work, and there are many divergent ways of going about it, all of which are controversial. It is commonly presumed, however, that what a sentence is conventionally used to assert does some pretty heavy lifting. A more precise explication of this line of thinking can be found in Soames (2002), which offers the following constraint on the semantic content of a sentence:

> C+: A proposition p is semantically expressed by s only if p is included in the information a competent speaker would assert and intend to convey by an assertive utterance of s in any context c in which s is used with its literal meaning by conversational participants who understand s, provided that (i) s is not used metaphorically, ironically, or sarcastically in c, and (ii) the presumption that the speaker intends to commit himself or herself to p is not defeated by a conversational implicature to the contrary. (Soames 2002: 60)

As Hawthorne notes, though on the right track, C+ is insufficiently general because it does not apply to sentences that we use in our own silent deliberations and practical reasoning. That said, an analogous constraint can be generated by invoking the notion of conscious acceptance:

> C++: A proposition p is semantically expressed by s only if, for any context [c] in which a competent speaker consciously accepts s, that speaker would thereby manifest a judgment that p, provided that s is not used metaphorically, ironically, or sarcastically in c. (Hawthorne 2004: 122)

Of course, as Hawthorne acknowledges, principles of this kind are obviously not incompatible with views on which the semantic contents of knowledge-attributing sentences are held to be something very different from what such sentences are typically used to convey, as proponents of such views can accept them while appealing to caveats (i) and (ii) in C+ and to the exception clause in C++. However, according to him, this presents a problem, as defenders of such accounts are left with the at best exceptionally difficult task of explaining just where they might get their semantic contents from instead. For

36 Whether it offers an equally good, or a better, explanation of these phenomena is again an exceptionally complicated issue and thus regrettably beyond the scope of this paper. A discussion of such questions should, however, be included as part of a complete case for epistemic fictionalism.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

presumably, such an explanation would have to include information about how sentences of the form "S knows p" behave in other contexts: otherwise, there would be no reason to suppose that they are being used non-literally in the contexts in question. However, it does not seem as if such information is forthcoming, as on the views under consideration, sentences of the form "S knows p" are typically used non-literally. Further, unlike expressions that are more plausibly used non-literally most of the time (such as "This line is taking forever"), sentences of the form "S knows p" do not have lexical constituents (relevant to the task at hand) that are held to have a sober use in some significant number of contexts. But, if the semantic content of "knows" rarely, if ever, plays a role in the propositions that are conveyed and believed when knowledge-attributing sentences are used, from where might their purported semantic contents derive? To the extent that one thinks that this question should be answered, and that there are no better places to look for an answer than those already canvassed by Hawthorne, one might be tempted to conclude at this point that views on which knowledge-talk is held to be largely non-literal are indefensible (Hawthorne 2004: 122–3).

I would like to question three assumptions about such views that Hawthorne appears to make. The first is that they must hold that "knows" does not have a sober use in some significant number of contexts. The second is that they must hold that the semantic content of "knows" rarely, if ever, plays a role in the propositions that are conveyed and believed when knowledge-attributing sentences are used. Finally, the third is that they must hold that the semantic content of an expression must somehow be determined by facts about speakers' uses of it in non-antagonistic contexts, rather than from facts about their uses of it in antagonistic contexts or facts that do not directly concern use at all.

Regarding the first assumption, it seems as if "knows" quite possibly does have a sober – and yet highly, if not impossibly, demanding[37] – use in a significant number of contexts on the fictionalist account proposed here. After all, we can observe on the basis of the discussion of retraction behavior in Section 2.7, and Camp's work on antagonistic contexts and their relevance in Section 3.2, that the literal, i.e., sober, use of "knows" might indeed be so demanding. We have also noted that skeptical arguments can be employed both within and outside of the philosophy classroom, and on a fairly regular basis – especially in antagonistic contexts – which gives us additional reason to suspect that sentences of the form "S knows p" may have more stringent truth-conditions than our rather permissive use of such sentences in non-antagonistic (i.e., fully, or highly, cooperative) contexts suggests. In connection with this, it is worth emphasizing that the examples provided at the outset of Section 3.2 exemplify just a few ways that skeptical arguments between people – which notably share the same basic structure as traditional skeptical arguments proposed in philosophy – might play out in ordinary conversations,

37 Indeed, perhaps so highly demanding that it is impossible to meet the satisfaction conditions that it mandates (which arguably rise to the level of infallibility). On whether knowledge might require infallibility, see, e.g., Unger (1975), Williamson (2000), and Dutant (2007, 2015). Some recent empirical work also suggests that a core folk concept of knowledge may be infallibilist, on the basis of an experimentally observed "certainty effect" in knowledge ascription (Machery 2016). Note, however, that infallibilism on its own may not require skepticism, or *vice versa*: some versions of infallibilism, such as that proposed in Williamson (2000), are meant to be antiskeptical, and it is well-known that many types of skeptical arguments can be ostensibly formulated so as to impugn degrees of epistemic support that fall far short of infallibility.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

especially in cases where agents are not disposed to interpret each other charitably. Nearly any everyday claim to know can be undermined in a similar fashion, as is evidenced by the fact that such arguments, if compelling, appear to suggest that we should be hesitant to attribute knowledge regarding propositions about the past, present, and future (as they clearly generalize beyond the specific situations to which they are meant to apply).[38] Further, while philosophical discussions about skepticism *per se* are not part of everyday English-language discourse, contradicting other people's perceived claims to know – overtly or covertly (even pedantically) – in antagonistic contexts is. Certain skeptical possibilities, such as the press making a mistake, a person lying, or an unforeseen disaster occurring are frequently admitted to be knowledge-imperiling when they are raised even if they are ultimately set aside. To see this, one need look no further than ordinary discussions regarding, say, politics or personal safety, in which conspiracy theories, calls for vigilance, and other types of far-fetched, but still possibly true, allegations or suspicions abound. In places such as these, as elsewhere, the old adage, "You never know," is a common refrain.

Moreover, it seems as if sentences of the form "S knows p" are no (or at least not a lot) worse off than "This line is taking forever" in this regard. For although "forever" can be used literally, it arguably is very rarely used in this way, unless one thinks that what was once a hyperbolic, non-literal use has been encoded into the semantics of the English language in such a way so as to make it a literal use. This, however, is not something that Hawthorne seems to think, as he appears willing to countenance the possibility that "This line is taking forever" is in fact typically used hyperbolically, and presumably, "forever," rather than any of the other terms in this sentence, would be to blame.

Hawthorne's second assumption is also problematic. For even if sentences of the form "S knows p" are not typically used to assert their semantic contents, it does not follow from this that the semantic content of "knows" rarely, if ever, plays a role in the propositions that are conveyed and believed when such sentences are used. To see more clearly why not, let us again take a look at what an epistemic fictionalist claims typically gets asserted when sentences of the form "S knows p" are uttered with declarative force. According to them, the content that a sentence of the form "S knows p" is customarily used to assert may usually be to some degree indeterminate over the set of propositions: {that S truly believes p, that S is justified in believing p, that S believes p as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to p, that S can assert p, that S can rely on p in practical reasoning, . . .}. Granted, none of the propositions mentioned here explicitly reference knowledge or knowing (although they tacitly might rely on some conception of them). But there may be others in the set that do. Possibilities include propositions along the lines of: that S's believing p is close enough to knowing p, that S's believing p resembles knowing p in some relevant respect, and that S can act as if S knows p, and so on. Of course, it would have to be the literal meaning of "knowing" or "knows" that they consider to be in play at this point, on pain of regress. So it seems as if, on their view, the literal meaning of "knows" actually can play a significant role in the propositions that are conveyed and believed when knowledge-attributing sentences are used.

Finally, Hawthorne's argument appears to rely on the unjustified and overly simplistic assumption that the semantic content of an expression must somehow be derived from

38 Cf. Hawthorne (2004).

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

facts about speakers' typical uses of it in non-antagonistic contexts, rather than from facts about their uses of it in antagonistic contexts or facts that do not directly concern use at all. It is, however, far from clear as to whether this is correct. For, as discussed above, antagonistic contexts are not only real but also sufficiently pervasive and illuminating so as to warrant inclusion in our theoretical deliberations – indeed, it may even be the case that we should look primarily to expressions' uses in antagonistic contexts in order to ascertain their semantic contents. Thus, not only is there no reason to assume that the uses of an expression that are responsible for fixing its semantic content must be its most common (or "majority") uses, but there is reason to think that we in fact should not take an egalitarian attitude toward all uses of "knows," and that we should prioritize certain uses – namely, those made in antagonistic contexts – over others, especially given that our retraction behavior is what it is.

Moreover, perhaps focusing solely on facts about use (whether in non-antagonistic or antagonistic contexts) is misguided from the get-go. Consider the possibility that semantic contents are grounded not in facts about speakers' uses of expressions considered alone or as such, but rather in facts about the broader communicative or cognitive purposes that speakers have that are associated with their uses of those expressions. For if not all uses of an expression have equal weight when it comes to determining its semantic content, something must determine their relative importance. Purposes may do this: after all, some uses of an expression might be more central to the achievement of our aims than others, and because of this, one might think that those uses should be accorded greater importance as far as determining its semantic content is concerned. However, if this is right, then although it could be the case that the use-facts could not change without the semantic facts changing, since any change in the use-facts would be due to a change in the purpose-facts (as, by hypothesis, purpose-facts determine use-facts – in other words, our purposes for using an expression will determine how we use it), it may be the case that the latter (the purpose-facts), rather than the former (the use-facts), are what ultimately determine the semantic content of an expression.[39]

If this hypothesis is correct (and perhaps even if not, as we may need to ask the following question anyway if we are to figure out how to weigh use-facts), in order to properly answer a question of the form "What is the semantic content of the expression e?" we would first need to be able to answer the question "What is the (primary) purpose of the expression e?"[40] In order to ascertain the semantic content of "knows," then, we

39 To see more clearly why this might be so, consider a prominent foundational theory of meaning that attempts to specify the facts in virtue of which expressions of natural languages come to have the semantic properties that they have. Perhaps the dominant approach to giving such an account in contemporary analytic philosophy of language is the so-called "mentalist" one. The most popular versions are the Gricean view, which explains meaning in terms of the communicative intentions of language users, and the view that the meanings of expressions are fixed by conventions which pair sentences with certain beliefs (cf. Speaks 2017). However, it is notoriously difficult to specify what, e.g., the relevant intentions of language users, or conventions in play, might be. One cannot give an account of either just by compiling a list of the many uses of expressions, especially since many of them will be non-literal. Rather, more is needed – including, at minimum, some explanation as to what the expressions in question are used for (in other words, an account of the relevant purpose-facts).

40 This is similar to following Edward Craig's (1990) advice to ask what the concept of knowledge does for us and treat our answers as clues about its application conditions (cf. McGrath 2015).

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

would first need to provide an answer to the question: What is the (primary) purpose of the expression "knows"? Or, in other words, what is knowledge-talk (principally) for?

One initially promising line of response to this question is that knowledge-talk plays a distinctive role in (conscious, deliberate) inquiry. Saying that it is not known whether p is true generally yields (an at least *pro tanto*) reason to start inquiring about whether p is true (for the first time or anew), as it introduces the idea that fruitful discussion about whether p is true is possible (provided that p isn't taken to be beyond our ken). Similarly, saying that it is known whether p is true generally yields (an at least *pro tanto*) reason to stop inquiring about whether p is true. For, it introduces the idea that fruitful discussion about whether p is true is not possible. We use knowledge-talk to challenge, but we also use it to give ourselves a place to rest. Although both of these functions of knowledge-talk are obviously crucial, the former arguably takes a certain sort of priority over the latter. For even though the ability to end inquiry about whether p is true (if only temporarily) is important to our proper functioning as epistemic agents, it seems as if the ability to begin inquiry about whether p is true is more fundamental. After all, one cannot end inquiry without first having begun it. Moreover, the main (if not the whole) point of ending inquiry about something seems to be to free up cognitive resources so that one might inquire about something else instead. If this much is granted, then it is possible that the primary purpose of knowledge-talk is to coordinate inquiry in a unique sort of way, either by inspiring us to begin inquiry, or end it (so that we can inquire about something else). And if the semantic content of an expression e is grounded in facts about the primary purpose of e, then the semantic content of "knows" might somehow be grounded in these facts about its distinctive role in coordinating inquiry (provided, of course, that its primary purpose is approximately as it is described here), rather than in facts about our uses of it *per se*. And if the more demanding "questioning-use" of "knows" is more fundamental (as opposed to the less demanding "resting-use"), then it may be better regarded as its literal meaning than the less fundamental "resting use".

Although I cannot flesh out and defend a more detailed proposal here, it should be clear enough from all of this that it is not in any way obvious that the semantic content of an expression must somehow be derived from facts about speakers' uses of it – rather, perhaps facts about purposes are more apt. To be sure, (at least some) facts regarding our uses of certain expressions are undeniably useful for determining (in the epistemic sense) just what the purpose of an expression is. After all, for the reasons just mentioned, at least some purpose-facts are likely to be revealed by use-facts. But it does not follow from this that facts about our uses of certain expressions determine (in the metaphysical sense) their semantic contents. Again, facts about purposes may well do that. And even if use-facts are what ultimately ground semantic content, in order to figure out how to weigh these, we might need to get into a discussion of purposes nonetheless.[41]

41 It should perhaps be noted that it has never been all that clear as to what the use of an expression is, exactly. Some theorists may employ the term very broadly, in which case the purposes of the speaker upon uttering the expression may be considered aspects of the use of the expression. But if one takes such a broad view of use the relevant problem for the fictionalist's proposal disappears. Thank you to Zoltán Szabó for pointing this out to me.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

3.6 Conclusion

I hope to have shown that the form of epistemic fictionalism outlined here is well-supported by several notable patterns of fictionalist argument. What's more, it can also be motivated by an additional argument concerning retraction behavior, and can weather some pretty tough objections.[42] This establishes it as a position worth pursuing further, especially since no other way of addressing the problems under discussion is even close to being widely accepted.[43]

REFERENCES

Blome-Tillman, M. 2009. 'Knowledge and Presuppositions.' *Mind*, 118: 241–94.
Blome-Tillman, M. 2014. *Knowledge and Presuppositions*. Oxford: Oxford University Press.
Brown, J. 2006. 'Contextualism and Warranted Assertibility Manoeuvres.' *Philosophical Studies* 130: 407–35.
Burgess, J. and Rosen, G. 1997. *A Subject with No Object*. Oxford: Oxford University Press.
Camp, E. 2006a. 'Contextualism, Metaphor, and What is Said.' *Mind and Language*, 21: 280–309.
—— 2006b. 'Metaphor and That Certain 'Je ne Sais Quoi".' *Philosophical Studies*, 129: 1–25.
—— 2008. 'Showing, Telling and Seeing. Metaphor and 'Poetic' Language.' *Baltic International Yearbook of Cognition, Logic and Communication*, 3: 1–24.
—— 2009. 'Two Varieties of Literary Imagination: Metaphor, Fiction, and Thought Experiments.' In H. Wettstein (ed.), *Midwest Studies in Philosophy: Poetry and Philosophy XXXIII*, pp. 107–30. Oxford: Wiley-Blackwell.
—— 2013. 'Insinuation, Inexplicitness, and the Conversational Record' (draft of 9/9/13 for Rutgers Semantics Workshop).
Chung, J. Forthcoming. 'Skeptical Arguments, Conceptual Metaphors, and Cross-Cultural Challenges.'
Cohen, S. 1999. 'Contextualism, Skepticism, and the Structure of Reasons.' *Philosophical Perspectives*, 13: 57–89.
—— 2000a. 'Contextualism and Skepticism.' *Philosophical Issues*, 10: 94–107.
—— 2000b. 'Replies.' *Noûs*, 34 (S1): 132–9.
Craig, E. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.
Davis, W. 2007. 'Knowledge Claims and Context: Loose Use.' *Philosophical Studies*, 132: 395–438.
—— 2015. 'Knowledge Claims and Context: Belief.' *Philosophical Studies*, 172: 399–432.

42 One question that this paper admittedly leaves wide open, but that is common in discussions of fictionalism is: What are the reasons for engaging in the fiction in question? While I cannot give a complete answer here, a few remarks stand to be instructive. First, some of the reasons for using knowledge-talk non-literally are likely to be much the same as the reasons for engaging in non-literal speech in general: to better grasp one thing by understanding and experiencing it in terms of another, to evoke certain associations and sentiments, to reduce cognitive load, and so on. Second, there are plausibly special reasons for using knowledge-talk non-literally in particular, such as: Using "knows" non-literally may encourage us to be curiously skeptical, rather than paralyzed by doubt, as doing so may permit us to act as if we are immune from error – but only to an extent – thereby allowing us to avoid the pitfalls of excessively entertaining doubts while at the same time shielding us from the dangers of dogmatism. For a more detailed response, please see Chung (Forthcoming). Thank you to an anonymous referee for suggesting that I address this.

43 I would like to thank Zoltán Szabó, Keith DeRose, Daniel Greco, Laurence Horn, Tamar Gendler, Jay Garfield, and Elisabeth Camp for their perceptive comments and invaluable guidance. I would also like to thank Mark Maxwell, Alex Worsnip, Zachary Gartenberg, and Gabriel Citron, as well as the many who participated in Yale's Semantics Reading Group, ELLMM City Working Group, WIP seminars, and philosophy retreats, for engaging and illuminating discussions.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

DeRose, K. 1995. 'Solving the Skeptical Problem.' *Philosophical Review*, 104: 1–52.
—— 2009. *The Case for Contextualism*. Oxford: Oxford University Press.
Dorr, C. 2005. 'There are no Abstract Objects.' https://files.nyu.edu/cd50/public/papers/NoAbstractObjects.pdf.
Dretske, F. 1970. 'Epistemic Operators.' *Journal of Philosophy*, 67: 1007–23.
Dutant, J. 2007. 'The Case for Infallibilism.' In C. Penco, M. Vignolo, V. Ottonelli and C. Amoretti (eds), *Proceedings of the 4th Latin Meeting in Analytic Philosophy*, pp. 59–84. Genoa: University of Genoa.
—— 2015. 'The Legend of the Justified True Belief Analysis.' *Philosophical Perspectives*, 29: 95–145.
Eklund, M. 2015. 'Fictionalism.' In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/win2015/entries/fictionalism/.
Fantl, J. and McGrath, M. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
Goldman, A. 1976. 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy*, 73: 771–91.
—— 1979. 'What Is Justified Belief?' In G. Pappas (ed.), *Justification and Knowledge*, pp. 1–25. Dordrecht: Reidel. Reprinted in Goldman, A. 1992. *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press.
Hawthorne, J. 2000. 'Reply to Cohen.' *Noûs*, 34: 117–20.
—— 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
Klein, P. 1981. *Certainty: A Refutation of Scepticism*. Minneapolis, MN: University of Minnesota Press.
—— 1995. 'Skepticism and Closure: Why the Evil Genius Argument Fails.' *Philosophical Topics*, 23: 213–36.
Lehrer, K. 2000. *Theory of Knowledge*, 2nd edition. Boulder, CO: Westview Press.
Lewis, D. 1996. 'Elusive Knowledge.' *Australasian Journal of Philosophy*, 74: 549–67.
MacFarlane, J. 2005. 'The Assessment Sensitivity of Knowledge Attributions.' In T.S. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology* 1, pp. 197–233. Oxford: Oxford University Press.
—— 2011. 'Relativism and Knowledge Attributions.' In D. Pritchard and S. Bernecker (eds), *The Routledge Companion to Epistemology*, pp. 536–44. New York, NY: Routledge.
Machery, E. 2016. 'Is There a Core Folk Epistemology?' In *International Conference on Ethno-Epistemology: Culture, Language, and Methodology*, 3–5 June 2016, Kanazawa, Japan.
McGrath, M. 2015. 'Two Purposes of Knowledge Attribution and the Contextualism Debate.' In D. Henderson and J. Greco (eds), *Epistemic Evaluation: Point and Purpose in Epistemology*, pp. 138–57. Oxford: Oxford University Press.
Nagel, T. 1986. *A View from Nowhere*. Oxford: Oxford University Press.
Nozick, R. 1981. *Philosophical Explanations*. Cambridge. MA: Harvard University Press.
Prades, J. 2000. 'Scepticism, Contextualism, and Closure.' *Philosophical Issues*, 10: 121–31.
Pritchard, D. 2005. 'Contextualism, Skepticism, and Warranted Assertibility Manoeuvres.' In J. Keim-Campbell, M. O'Rourke and H. Silverstein (eds), *Knowledge and Skepticism*, pp. 85–104. Cambridge, MA: MIT Press.
Pryor, J. 2000. 'The Skeptic and the Dogmatist.' *Noûs*, 34: 517–49.
Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
Pynn, G. 2015. 'Pragmatic Contextualism.' *Metaphilosophy*, 46: 26–51.
Rysiew, P. 2001. 'The Context-Sensitivity of Knowledge Attributions.' *Noûs*, 35: 477–514.
—— 2007. 'Speaking of Knowing.' *Noûs*, 41: 627–62.
Schaffer, J. 2004a. 'From Contextualism to Contrastivism.' *Philosophical Studies*, 119: 73–103.
—— 2004b. 'Skepticism, Contextualism, and Discrimination.' *Philosophy and Phenomenological Research*, 69: 138–55.
—— **and Szabo, Z.** 2014. 'Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions.' *Philosophical Studies*, 2: 1–53.
Soames, S. 2002. *Beyond Rigidity*. Oxford: Oxford University Press.
Sosa, E. 1994. 'Philosophical Scepticism and Epistemic Circularity.' *Aristotelian Society Suppl.* 68, 263–90.
Speaks, J. (2017) 'Theories of Meaning.' In E.N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/spr2017/entries/meaning/.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press

Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
Stine, G. 1976. 'Skepticism, Relevant Alternatives, and Deductive Closure.' *Philosophical Studies*, 29: 249–61.
Stoutenburg, G. Forthcoming. 'Strict Moderate Invariantism and Knowledge-Denials.' *Philosophical Studies*.
Stroud, B. 1984. *The Significance of Philosophical Skepticism*. Oxford: Oxford University Press.
—— 2011. 'Critical Study of Mark Eli Kalderon (ed.) *Fictionalism in Metaphysics*.' *Noûs*, 45: 375–85.
Szabó, Z. 2001. 'Fictionalism and Moore's Paradox.' *Canadian Journal of Philosophy*, 31: 293–307.
Unger, P. 1971. 'A Defense of Skepticism.' *Philosophical Review*, 80: 198–218.
—— 1975. *Ignorance*. Oxford: Oxford University Press.
—— 1984. *Philosophical Relativity*. Minneapolis, MN: University of Minnesota Press.
—— 1986. 'The Cone Model of Knowledge.' *Philosophical Topics*, 14: 125–78.
Vogel, J. 1990a. 'Cartesian Skepticism and Inference to the Best Explanation.' *Journal of Philosophy*, 87: 658–66.
—— 1990b. 'Are there Counterexamples to the Closure Principle?' In M. Roth and G. Ross (eds), *Doubting: Contemporary Perspectives on Skepticism*, pp. 13–27. Dordrecht: Kluwer.
Walton, K. 1993. 'Metaphor and Prop Oriented Make-Believe.' *European Journal of Philosophy*, 1: 39–57.
Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
Wright, C. 2004. 'Warrant for Nothing (and Foundations for Free)?' *Aristotelian Society Suppl.*, 78: 167–212.
Yablo, S. 1998. 'Does Ontology Rest on a Mistake?' *Aristotelian Society Suppl.*, 72: 229–6.
—— 2000. 'A Paradox of Existence.' www.mit.edu/~yablo/apex.html.
—— 2001. 'Go Figure: A Path Through Fictionalism.' *Midwest Studies in Philosophy*, 25: 72–102.

JULIANNE CHUNG is Assistant Professor of Philosophy at the University of Louisville. She is especially intrigued by questions having to do with skepticism, fictionalism, and metaphor, as well as how cross-cultural philosophy can shed new light on them. She is also interested in whether (and how) works of art – and, in particular, literature – can have cognitive content, and what the philosophical upshots of this might be. She is presently associate editor of *Oxford Studies in Epistemology*, and has been recently nominated to the American Philosophical Association's committee on Asian and Asian-American Philosophers and Philosophies.

https://doi.org/10.1017/epi.2017.7 Published online by Cambridge University Press