WILEY

Assertoric content, responsibility, and metasemantics

Andrew Peet (ORCID)

The School of Philosophy, Religion and History of Science, University of Leeds, Leeds, UK

Correspondence
Andrew Peet, The School of Philosophy, Religion and History of Science, University of Leeds, Leeds LS2 9JT, UK.
Email: andrewpeet123@gmail.com

I argue that assertoric content functions as a means for us to track the responsibilities undertaken by communicators, and that distinctively assertoric commitments are distinguished by being generated directly in virtue of the words the speaker uses. This raises two questions: (a) Why are speakers responsible for the content thus generated? (b) Why is it important for us to distinguish between commitments in terms of their manner of generation? I answer the first question by developing a novel responsibility based metasemantics. I answer the second by reference to the conflicting pressures governing the resources we have available for appraising speech.

KEYWORDS
assertion, commitment, communication, content, context sensitivity, metasemantics

1 | INTRODUCTION

That we assert things to one another, and that our doing so is central to our linguistic practice, seems beyond question.[1] When we assert, there is typically something we can be said to have asserted. This is what we might think of as the truth conditional content of our assertion. It is

---

[1]Although see Cappelen (2011, 2019) for a dissenting view. It is an assumption of this article that we have a reasonably clear pretheoretic notion of assertion and assertoric content and that this provides the basis for philosophical discussions of, for example, the norms of assertion. The arguments and conclusions of the paper should be taken as conditional on this assumption.

....

the content relative to which we define norms of assertion. That is, it is the content that must be true (or perhaps known, believed, or warranted for the speaker) if an assertion is to deemed normatively acceptable. It is hard to make sense of the notion of assertion without the notion of assertoric content. Yet, as I will illustrate in the earlier portions of this article, it is not immediately clear what communicative function assertoric content actually plays.

Assertoric content must, I argue, be distinguished from both intended content and the content determined by standing meaning. Cases of error, in which a speaker accidentally asserts p despite their intention to assert some alternative proposition q, force us to distinguish assertoric content from intended meaning. Cases of modulation drive a wedge between assertoric content and standing meaning. Yet, once separated from standing meaning and intended meaning, it is not clear what role assertoric content could play. The aim of communication is for interlocutors to coordinate on intended meanings. Communication breaks down when the audience fails to recover the speaker's intended meaning. And communication is facilitated by knowledge of standing meaning. That is, knowledge of standing meaning, together with contextual knowledge, typically suffices for the audience to recover the speaker's intended meaning, and thus for communication to succeed. It is not clear what role assertoric content could play in this process. If assertoric content plays no role in this process, it would seem to be superfluous. Yet, given the centrality of assertion to our thought and talk about language use, assertoric content cannot be superfluous.

In order to resolve this puzzle, I argue, we must identify an alternative function for assertoric content. The basic idea is this: Assertoric content functions as a means for us to track the responsibilities undertaken communicators when they speak. However, this, by itself, is not very illuminating. There any many things we take responsibility for when we speak, and many ways in which we undertake such responsibilities. Not all of the responsibilities we undertake when we assert will correspond to the intuitive notion of assertiotic content.

I suggest that assertoric commitments are distinguished by their mode of generation: They obtain directly (either compositionally or via bridge principles) in virtue of the words, the speaker uses and their manner of combination. But this raises two further questions: First, why are speakers responsible for the content thus generated? And second, why is it important for us to distinguish between communicative commitments in terms of the manner in which they are generated? Put another way: What justifies the central place assertoric commitments seem to occupy in our practices of normatively appraising speech if they only differ from other communicative commitments in their manner of generation?

My primary focus will be on the first question: I argue that a plausible metasemantic theory must be able to make sense of the fact that speakers are committed to the assertoric contents of their utterances. Some metasemantic theories are better equipped to do this than others. I present a metasemantic theory that is particularly well equipped to do so: The value a term receives in context corresponds to the use it is most fitting (i.e., there is the most reason) to hold the speaker to in light of their utterance. I turn to the second question in the conclusion.

2 | ASSERTORIC CONTENT

As I noted above, when a speaker asserts there will typically be some proposition that they assert. The truth or falsity of their assertion will depend on whether or not this proposition is true. Likewise, their conformity to the norm of assertion will depend on their relation to this

proposition (i.e., Do they know the proposition? Do they believe it?). This is what I have in mind when I speak of assertoric content.

2.1 | Assertoric content and semantic content

Assertoric content will be largely determined by the semantic values of the words used by a speaker in context. Indeed, assertoric content and semantic content are often equated. However, their equation is controversial. First, even theorists who reject traditional externalist semantics (such as, Chomsky, 1995; and Pietroski, 2003) can endorse a notion of assertoric content. For example, Pietroski tells us that:

> We should be sceptical of the idea that a theory of meaning for a natural language will have theorems that specify the truth-conditions of all the declarative sentences of that language. The successes in semantics suggest that the theoretical action lies elsewhere; semantics is concerned with "internalist" features of linguistic expressions, rather than truth per se. The fact that (an utterance of) a sentence has a certain truth-condition is typically an interaction effect whose determinants include (i) intrinsic properties of the sentence that we can isolate and theorize about, and (ii) a host of facts less amenable to theorizing, like facts about how "reasonable" speakers would use the sentence. (Pietroski, 2003, p. 1)

For Pietroski assertoric content is determined by semantic meaning (internalistically conceived), and other factors which he takes to be less amenable to systematic theorizing.[2]

Second, semantic minimalists such as Borg (2004, 2012), and Cappelen and Lepore (2005) countenance a minimal propositional and context insensitive form of semantic content that rarely (if ever) corresponds to assertoric content, which tends to be richer, more informative, and also highly context sensitive (although Borg, 2019, argues that minimal content does coincide with assertoric content in certain special contexts).

Finally, even adherents of more traditional approaches to semantics have questioned the identification of semantic content with assertoric content. For example, Dummett (1973), Evans (1979), Lewis (1980), Stanley (2002), Ninan (2010), Rabern (2012, 2017), and Yalcin (2014) present a number of conceptual and empirical arguments against the identification.[3] Regardless of whether or not semantic content is to be identified with assertoric content, there is a close relationship between the two. The contextual values of the words used (i.e., the values the terms are assigned on an occasion of use), whether they be systematically context sensitive terms like "tall," or context insensitive terms (which might still be modulated or used loosely) such as "raw," largely determine (either compositionally, or via certain bridge principles) assertoric content. This will be crucial in what follows.

---

[2]It might be worried that the project I am engaged in here is in conflict with the Chomsky/Pietroski view. After all, I will be giving an account of the contextual values (e.g., referents) of terms on occasions of use. And this seems to be precisely what Chomsky and Pietroski are sceptical of. I will return to this issue in the conclusion when I argue that the approach advocated here actually captures the sense in which semantics is beyond the scope of systematic empirical investigation.

[3]See Stonjić (2017) for a reply.

2.2 | Assertoric content and intended content

Assertoric content must be distinguished from intended content. Intended content is the proposition the speaker aims to directly communicate with their assertion. That is, its the primary proposition the audience must recover if the speaker's communicative intention is to be satisfied. A core assumption of this article is that communicative success ultimately turns on the recovery of intended content. That is, in typical cases of communication, the speaker aims for the hearer to recover the content they intend to communicate, and the hearer aims to recover the content the speaker intends. Coordination on intended content is the core aim of communicative interactions.[4]

When all goes well assertoric content and intended content will coincide. However, the two can come apart, especially in cases of speaker error. Examples abound from the literature on the metasemantics of indexicals, demonstratives, and gradable adjectives. Such cases are controversial. However, these controversies need not concern us since we can present analogous cases that do not turn on the values assigned to context sensitive terms:

> WHALES: Smith goes on a whale watching trip. He sees a humpback whale, and recognises it as such. When he returns to shore he is asked whether he saw anything. In a moment of confusion, he misremembers the name for humpbacks and, in a sincere attempt to express his belief about the whale he saw, he states "yes, I saw a sperm whale." He intended to communicate that he saw a humpback whale. In some sense this is what *he* meant. However, this is not what he asserted.

Smith asserted something false: that he saw a sperm whale. It would stretch credulity to claim otherwise. Thus, we have a clear case in which intended content and assertoric content come apart. Of course, this does not show that speaker intentions play no role in *determining* assertoric content. It could be, for example, that standing meanings constrain what certain terms can be used to mean, with speaker intentions fixing the values of context sensitive terms, or resolving indeterminacy. Alternatively, as I shall suggest later, it may be that speaker intentions are one among a number of factors which must be weighed to determine assertoric content. The considerations presented so far merely show that intended content and assertoric content do not always coincide.

However, if intended content and assertoric content are not to be identified, and communicative success turns on the recovery of intended content rather than assertoric content, then it becomes unclear what role assertoric content might play in communication.

2.3 | Assertoric content and standing meaning

When considering cases like WHALES, it might be thought that we can identify assertoric content with the content determined by the standing meaning of the sentence used (call this level of content "standing content"). After all, in WHALES the proposition Smith asserted was that

---

[4]In reality, the situation may be more complex than this. It may be that precise coordination on intended content is rare, and that interlocutors need only coordinate on sufficiently similar propositions. However, even if this is the case, the proposition the hearer recovers will have to be sufficiently similar to the speaker's intended proposition in order for communication to succeed (see Peet, 2019, for a view along these lines).

corresponding to the standing meaning of "sperm whale," rather than the meaning he had in mind.

Moreover, if assertoric content was to be identified with standing content the function of assertoric content would be less puzzling: Standing meaning facilitates the communication of intended content. The mechanisms by which it does so are, of course, controversial. But on a simple model they can be seen as providing a starting point for the audience's subpersonal inferential processing, the output of which is a representation of the intended content. It is standing meaning which will be modulated (or, perhaps, filled in or precisified) by the audience in an attempt to recover intended content.

Unfortunately, assertoric content cannot be identified with standing content. There are several reasons for this. First, as previously noted, assertoric content typically coincides with intended content. Assertoric and intended content usually only come apart when the speaker has committed some error. Standing content corresponds to standing meaning, and standing meaning is the primary input to interpretation; that which is modulated to reach a representation of the intended content. Thus, at least if context sensitivity and modulation are the norm (which I assume they are), standing content and intended content will not typically coincide.

Relatedly, many theorists argue that standing meaning will not typically determine a complete proposition (Bach, 1994; Carston, 2002; Recanati, 2004; Sperber & Wilson, 1986). If this is correct, then standing content, unlike assertoric content, will typically be highly indeterminate.

Finally, there are intuitive counterexamples to the identification of standing content with assertoric content:

> STEAK: Martha and Bill are eating out at a restaurant. Martha has ordered a well-done steak. When it arrives, she cuts into it only to find that it has been served rare: It is seared on the outside, and only very lightly cooked in the middle. The steak is in clear view of Bill. Martha looks to Bill and says, "my steak is raw."[5]

The term "raw" in English means "uncooked." Thus, the standing content for Martha's utterance (once we have adjusted for the indexicality of "my steak") is that her steak is uncooked. However, this is not what Martha has asserted. She has not asserted anything false, nor has she violated the norms of assertion. What she really asserted is that her steak is undercooked. In this context, the term "raw" is modulated and takes on a wider meaning. The assertoric content of her utterance corresponds not to its standing content, but rather the content determined by the modulated meanings of the terms used.[6]

---

[5]Similar cases abound. Particularly clear is Nunberg's (1979) case of the waitress who asserts "the ham sandwich left without paying."

[6]This argument depends on our thinking of "standing meaning" as the meaning of a sentence in a relatively stable conventional public language (e.g., English or French), or a relatively stable idiolect. I take this usage to be standard. There are alternative conceptions of standing meaning according to which speakers typically operate with dynamic microlanguages developed and altered on the fly (Armstrong, 2016; Ludlow, 2014). Proponents of such approaches could plausibly claim that, in the microlanguage spoken by Martha and Bill, "raw" means "undercooked." In this case, the content of Martha's assertion will correspond to what we might call the "dynamic standing content" of her utterance. This will not resolve the problem, however, for we are left asking what the function of dynamic standing content is. If it functions as a primary input to interpretation then, assuming the ubiquity of context sensitivity it will, unlike assertoric content, typically fail to coincide with intended content. If we deny standing content this function then we are free to identify it with assertoric content, but we are left with the mystery as to its function.

2.4 | The function of assertoric content

Assertoric content should be distinguished from both from intended content, and standing content. Intended content and standing content both have clear communicative functions: Intended content is what interlocutors aim to coordinate on. Communicative success or failure will depend on whether or not the audience recovers the speaker's intended content (or, perhaps, a sufficiently similar proposition). And standing content corresponds to standing meaning, which facilitates the communication of intended content.

Unlike intended and standing content, assertoric content lacks a clear communicative function. Coordination on assertoric content is not required for successful communication. It is true that, typically, recovery of an utterance's assertoric content will suffice for communicative success. But this is only because assertoric content typically coincides with intended content. When the two come apart, recovery of assertoric content will not suffice for communicative success. Moreover, it is not clear what role assertoric content could play in facilitating communication. Assertoric content typically coincides with the output of interpretation, not with the input to interpretation, or the output of some intermediate stage in the interpretative process. So why, then, do we have a concept of assertoric content? And why does it appear to be so important?[7]

The function I propose is simple, and has been proposed by others: Assertoric content functions as a means for us to hold speakers responsible for their utterances. In general, our speech has an impact on the world. It changes beliefs, hurts feelings, and alters the normative relations we stand in to others. Sometimes this impact is positive, and sometimes it is negative. Either way, it is imperative that we have a resource with which to clearly and unambiguously hold one-another responsible (or creditworthy) for the potential impact of our speech. This is why assertoric content is so important: It tracks the content speakers make themselves responsible for through their speech (see Perry, 2006; Borg, forthcoming, for similar approaches to assertoric content, and for examples of attempts to characterize the speech act of assertion in terms of the type of the type of commitment assertions generate see Brandom, 1983, 1994; MacFarlane, 2011; Shapiro, 2020).[8]

As just noted, I am not the first to suggest this function for assertoric content. However, current attempts to elucidate this connection run into a number of problems, and leave some important questions unanswered. Perry (2006) does not go into much detail, but Borg (2019) puts the notion of what she calls "linguistic liability" to extensive use. She argues that we can

---

[7]It could, of course, be held that assertoric content, as distinct from intended and standing content, is superfluous (Richard Kimberly Heck, 2014, comes close to suggesting as much). However, I believe we should resist this hypothesis, adopting it only as a last resort. After all, assertions certainly *seem* to have truth conditions. There certainly seems to be a proposition asserted in cases like WHALES and STEAK. Surely there is some good reason for us to have these intuitions. Moreover, we have established practices which revolve around assertoric content. For example, the norms of assertion play a key role in our practices of normatively appraising speech. These norms are defined in terms of assertoric content, not intended or standing content. For these reasons, I think our first response should be to identify an alternative role for assertoric content. That is, we must ask what role assertoric content could play if it plays no role in communication.

[8]Perry (2006) claims that the term "what is said" is a "forensic concept," having to do with the message a speaker is responsible for conveying. I have avoided the term "what is said" here and opted to focus on "assertoric content" because I worry that "what is said" is used in different ways by different theorists, and that this leads to potential confusion. However, the general claim is similar. Likewise, Borg (2019) argues that both minimal content and explicature track distinct levels of content with different functional roles: minimal content tracks content a speaker is "strictly liable" for, and explicature tracks the content a speaker is "conversationally liable" for. She takes assertoric content to sometimes correspond to minimal content, and sometimes correspond to explicature.

carve out roles for minimal content, explicature, implicature, and assertoric content in terms of linguistic liability. She takes "strict" linguistic liability to track minimal content, she takes implicature and explicature to correspond to different levels of what she calls "conversational liability," and she takes assertoric content to correspond to either explicature or minimal content depending on the type of liability relevant in context. We might say that a speaker is "assertorically liable" when they produce an utterance with a minimal content p, and they are in a context in which strict liability is key (e.g., a legal context), or else they are highly conversationally liable for p, and they are in a context in which conversational liability is key.

Although I am broadly sympathetic to Borg's approach, there are some important differences between Borg's development of the responsibility view and my own. First, and most importantly, although Borg employs the notion of linguistic liability, she does not provide an account of linguistic liability. A large part of what I do in the following sections consists in giving an account of assertoric liability. That is, I suggest an account of the conditions under which a speaker is assertorically liable for a particular proposition. Moreover, in doing so, I go beyond providing an account of the function of assertoric content: I provide an account of the *nature* of assertoric content, and I provide an integrated metasemantics of context sensitive terms.

A second key difference concerns the way in which explicatures are distinguished from implicatures. Borg takes implicature and explicature to correspond to different levels of linguistic liability. So, for Borg, one is always more liable for what is said than what is implicated. As a rough generalization this is probably correct. However, it does seem to admit of exceptions. Consider the following:

> LETTER WRITER: Paul Grice famously gave the following example of implicature:
> A is writing a testimonial about a pupil who is a candidate for a philosophy job, and his letter reads as follows: "Dear Sir, Mr. X's command of English is excellent, and his attendance at tutorials has been regular. Yours, etc." (Grice, 1989, p. 33)

In this case, A has implied that Mr X is not a good candidate. Now suppose that Bill and Bob are both philosophers of language at different universities. They both work on implicature, and they each have a particular interest in Grice. Bill is writing a letter of reference for his rather unimpressive student Alex, and he knows that Bob (and only Bob) will be the one to read the letter. He also knows that Grice's letter writer example is salient common knowledge between the two of them. So, in his letter, he simply writes "Alex's command of English is excellent, and his attendance at tutorials has been regular."

The implication is clear: Alex is not an impressive student. Given the mutual knowledge between them, Bill is strongly committed to this. Indeed, I would suggest that in these circumstances, given their mutual knowledge of the Gricean example, his commitment to the claim that Alex is a poor student is actually stronger than his commitment to the claim that Alex has a solid grasp of English. Nonetheless, all Bill has said (and all Bill has asserted) is that Alex has an excellent grasp of English, and attends tutorials regularly. He has merely implied that Alex is a poor student. If this is correct, then we cannot make sense of the implicature/explicature distinction in terms of different grades of conversational liability.[9]

---

[9]These are not the only ways in which I depart from Borg, only the most salient. Borg also argues that the notion of strict liability is closely tied to the concept of lying. That is, to tell a lie, one must be strictly liable for something one believes to be false. I am skeptical as to the very notion of strict liability and minimal content. But even if I were to

Although he does not explicitly frame his discussion in terms of the function of assertoric content, Stainton (2016) also argues that assertoric content can be distinguished from merely implied content by virtue of the of commitment involved. However, unlike Borg, Stainton takes the commitments to differ not in terms of strength, but rather in terms of type. Stainton suggests that assertion generates a distinctive type of commitment characterized by a distinctive form of normative failure. Commitment accounts of assertion (such as those surveyed in Shapiro, 2020) can also be seen as characterizing the act of assertion in terms of the *type* of commitment one undertakes when one asserts. And there have been recent attempts to makes sense of the lying/misleading distinction in terms of the distinctive *type* of commitments involved in lying as well (e.g., Viebahn, forthcoming). I will focus on Stainton, but indicate the ways in which my worries generalize.

Stainton suggests that, just as promise breaking and adultery are distinctive types of normative failure that could not exist independently of the specific commitments generated by promising or marriage, lying is a form of normative failure that could not exist without the act of assertion. It is widely accepted that lying requires asserting. However, it is not obvious that lying is a distinct type of normative failing. Rather, I suspect that lying is one species of a more general normative failing: attempting to communicate something one believes to be false. Of course, the introduction of assertion as a means of communication introduces the possibility of failing in this way. But that does not mean it introduces a distinct form of normative failure. When the practice of sarcasm was first introduced, did this introduce a new form of normative failure: attempting to communicate something one believes to be false via sarcasm? Surely not! What emerged was just a new way to go about failing in a familiar manner. So, why should we consider the act of attempting to communicate something one believes to be false via assertion to be a distinctive normative failure? We do have a label for attempting to communicate something one believes to be false via assertion, and we do not have a label for the act of doing so via sarcasm. But giving a normative failing a label is not enough to make it distinctive!

It might be argued that there are independent grounds for thinking that assertoric commitment differs in kind from the sort of commitment we undertake when we merely imply. For example, Fricker (2012) has suggested that while we are able to maintain plausible deniability about the contents of implicatures, we are not able to maintain plausible deniability about the contents of assertions. However, there are reasons to doubt this: First, as Hawthorne (2012) and Peet (2015) have argued, the difference here is really only one of degree. In general, we tend to have a greater degree of plausible deniability with respect to implied contents. However, since language is rife with context sensitivity, and even the recovery of assertoric content requires significant inference on the audience's behalf, there will often be ways for a sufficiently imaginative speaker to maintain some level of plausible deniability about asserted contents. And, as we saw above, there are cases in which we are not able to maintain plausible deniability about the contents of implicatures (for similar points see Camp, 2018; Viebahn, 2017, 2020, forthcoming).

---

countenance these notions, I believe that strict liability is not necessary for one to have lied. Consider the following example:

> SANDWICH LIE: Mark is a waiter at a cafe. He is also a thief. When a customer who has ordered a ham sandwich pays and leaves (taking the sandwich with him) Mark pockets the money and tells the manager "the ham sandwich left without paying" (cf. Nunberg, 1979).

Mark has clearly lied. Moreover, his lie has nothing to do with the ham sandwich itself leaving the restaurant without paying. His lie concerns the action of the customer. Indeed, since the customer took the ham sandwich with him, and the ham sandwich did not pay, the minimal content was actually true.

Moreover, even if we were able to maintain plausible deniability about implied but not assertoric content, it is not clear that this would render the commitments different in kind. Finally, is not clear in what sense the commitments undertaken in LETTER WRITER really differ from the commitments typical of assertion. The letter writer will incur negative reactive attitudes and will be held to the content of their implicature if it is false just as an asserter would. This worry generalizes beyond Stainton's approach to commitment accounts of assertion in general.[10]

So, I am not optimistic about the prospect of distinguishing between assertoric and implied content by reference to the *type* of commitment we bear to each. Nor am I optimistic about doing so in terms of the *degree* of commitment involved. This raises a question: In what sense *do* the commitments generated by assertion and implicature differ? The key difference, I believe, is the mode of generation. Put simply, assertoric content, unlike, say, implied content, is determined directly, either compositionally or via bridge principles by the contextual values assigned to the worlds used.

But this just raises new questions: First, why should we be held responsible for the content determined this way? That is, why is the content determined by the contextual values of the words used in context something that speakers make themselves responsible for? Why is it important for us to distinguish between communicative commitments in terms of the manner in which they are generated? And what justifies the central place assertoric commitments seem to occupy in our practices of normatively appraising speech if they only differ from other communicative commitments in their manner of generation? We will start with the first question. I will turn to the second question in the conclusion.

3 | METASEMANTICS, USE, AND RESPONSIBILITY

I closed the previous section by suggesting that assertoric commitments are those that are determined directly, either compositionally or via bridge principles by the contextual values assigned to the worlds used. However, I also noted that this raises a question: Why, given the determinants of contextual value, is the content so generated something we should be held responsible for? This question will be easier to answer on some metasemantic views than others.

Consider the following toy metasemantics: The contextual value assigned to a term on an occasion of use corresponds to the use of that term most commonly intended by speakers in the relevant community. This metasemantic theory obviously has a number of problems. One problem is that it is not clear why it would make sense to hold speakers responsible for the content so generated. The theory can be rejected on these (and many other) grounds. Other metasemantic theories do better. Consider, for example, the idealized interpreter approach: The view that the value a context-sensitive term receives on an occasion of use is the value a suitably idealized interpreter would assign.[11] The idealized interpreter view fits naturally with our

---

[10]The problems discussed here also generalizes to Viebahn's (2020) commitment-based approach to the lying–misleading distinction. First, if the letter writer knew that the student was an excellent candidate then they would still have merely mislead, not lied. Yet they seem robustly committed to the claim that the student is a poor candidate. Second, Viebahn characterises the type of commitment liars undertake in terms of their ability to maintain coherent deniability regarding the content of their speech act. But as Peet (2015) points out, we are sometimes able to maintain a degree coherent deniability regarding the content of our assertions as well. We would presumably still want to hold that speakers have lied when they intentionally assert falsehoods whilst retaining a degree of deniability.

[11]See Wettstein (1984), and Romdenh-Romluc (2002, 2006). Additionally, King (2013, 2014a, 2014b), and Armstrong (2016) embrace hybrid views according to which, if a context sensitive term is to take on a particular value, the speaker must intend that it do so, and it must be such that a suitably idealized interpreter would assign that value.

hypothesis about assertoric commitments. After all, if a reasonable hearer would take a speaker to be using their words in a particular way then, surely, the speaker should be held to this use: It is this use that will determine the likely impact of their speech. Unfortunately, despite the initial appeal, idealized interpreter approaches are untenable. In the next section, I will explain why. Following this, I will outline a superior alternative.

3.1 | The idealized interpreter approach[12]

The idealized interpreter approach is subject to a number of counterexamples. While most of these counterexamples can be overcome with suitable idealizations, the cases will often pull us in conflicting directions. This renders the idealized interpreter approach untenable. Idealized interpreter approaches will hold that something like the following is true:

> IDEAL: The value a context-sensitive term receives on an occasion of use is whatever value an interpreter meeting the following conditions—would take the speaker to intend.

Different approaches will fill in the blank in different ways. The problem is that no single way of filling in this blank is satisfactory. Consider the following three ways of filling in the gap:

> GENERIC: The interpreter will be a generic hearer perhaps idealized so as to meet certain conditions of competence and contextual knowledge.

> ACTUAL: The interpreter will be a version of the speaker's actual audience, perhaps idealized so as to meet certain conditions of competence and contextual knowledge.

> EXPECTED: The interpreter will be an idealized version of the speaker's expected audience (i.e., an audience corresponding to the speaker's expectations of their audience) perhaps idealized so as to meet certain conditions of competence and contextual knowledge.

These are the obvious ways of developing an idealized interpreter approach. Let us start with GENERIC. Consider the following:

> X-RAY: Christina is teaching a course to advanced medical students. She is showing the students an x-ray image. This image would, to any normal viewer, be utterly unintelligible. However, the medical students possess specialist knowledge, and can interpret the x-ray easily. She points to an area with what, to the normal observer, would appear to be a number of grey blobs. She says, intending to refer to the blob in the centre, "that is a rupture." Her students, due to their skill in reading x-rays, and their prior medical knowledge, correctly recognise the blob she is referring to. They take her to be saying, of the blob, that it is (or, represents) a rupture.

---

[12]The argument presented here is brief, as Nowak and Michaelson (2021) have recently presented a similar case against the idealized interpreter approach.

A generic competent hearer would not be able to assign a determinate referential intention to Christina. They would have no way of telling which blob she intended to refer to. Thus, GENERIC seemingly predicts that Christina fails to refer to any particular blob. This is the wrong prediction. Christina said that the central blob is a rupture.

Of course, GENERIC could be further filled in so that it gets the correct result. We might maintain that our generic hearer must know the conversational common ground (King, 2013, 2014a, 2014b, adopts this approach), or we might maintain that they must meet certain of the speaker's expectations. But such idealizations push us away from GENERIC and toward versions of ACTUAL or EXPECTED.

Next consider ACTUAL: It is clear that ACTUAL gets the correct result with respect to X-RAY. Christina's actual audience, due to their specialist knowledge, is able to correctly recover the intended content of her utterance. Unfortunately, ACTUAL is rendered untenable by the following case:

> SPEAKS'S LECTERN: I teach Philosophy 101 in a large auditorium which darkens during the lecture so that the students can better see the slides; in fact, though, it becomes a bit darker than it needs to, to the point where I can't see the students during the lecture. The students have figured this out, and now, very quietly, exit the room minutes after the lights go down, and return minutes before the lights go back up. In the interim, I'm speaking to an empty room ... During the lecture, I might use plenty of demonstratives; and it seems to me very clear that many of them might well have semantic values. If, pointing clearly and carefully at the lectern, I say "that lectern ...," it seems clear that the semantic value of "that lectern" is, just as it would be had the students not left, the lectern. (Speaks, 2016, p. 312)

Since the lecturer has no audience, there is no object that the actual audience (or an idealized version thereof) would take him to be referring to. Thus, ACTUAL seemingly predicts that the lecturer's use of "that lectern" fails to receive a contextual value. Once again, this is the wrong result. The lecturer referred to the lectern. The truth or falsity of his utterance will depend on whether or not the lectern has whatever property he predicated of it.

EXPECTED does better here. The lecturer takes himself to be addressing a group of normal students. Any audience satisfying the lecturer's expectations would successfully recognize his intention. Unfortunately, EXPECTED is also subject to counterexample. Consider the following:

> FRENCH INVASION: Fernando is taking part in a re-enactment of the Norman Conquest. The re-enactment is to be staged in a manor house. Fernando is to play the part of Harold's messenger, who announces to the King that the French are invading England. At the appointed time, Fernando bursts into the main hall and announces to the assembled guests,
>
> (2) Now the French are invading England!
>
> However, Fernando has confused the date of the re-enactment. He has arrived a week early and made his announcement to wedding guests who are staying at the manor, not other people involved in the re-enactment of the Norman Conquest.

> The wedding guests are all panic-stricken, as they think France has just declared war on England. (Romdenh-Romluc, 2002, p. 37)

This case is more fanciful than the others. But if we imagine it taking place in a context of extreme tension between the UK and France (perhaps the Brexit negotiations have really gone south) it becomes plausible that Fernando has accidentality stated that the French are *currently* invading England. Yet EXPECTED predicts otherwise: Fernando takes his audience to be aware of the pretence. An audience meeting his expectations would take his use of "now" to refer to 1,066. This pushes us back toward versions of ACTUAL or GENERIC. But we have seen that such approaches are subject to counter example. It is possible that further patches could be applied to the idealized interpreter approach (although Nowak & Michaelson, 2021, argue that this is unlikely). However, I think this problem is sufficiently pressing that it is worth looking for an alternative.

3.2 | The fittingness approach

We are looking for a metasemantic theory that makes good sense of the fact that speakers should be held responsible for the contents directly generated, either compositionally or via bridge principles, by the contextual values assigned to the worlds used. The idealized interpreter approach was promising in this regard. However, it appears untenable on closer inspection. It is time to consider an alternative.

Just as we can be held to particular claims, we can be held to particular word uses. Returning to our earlier examples, there is a clear sense in which Martha should not be held to the standard use of "raw," and Smith, despite his intentions, should be held to the standard use of "sperm whale." My proposal, roughly, is that the contextual value a word receives on an occasion simply corresponds to the use of that word it is most fitting to hold the speaker to. Assertoric content, being a function of contextual values, will thus correspond to the proposition it is most fitting to hold the speaker to in light of the words used. Thus, the fittingness view provides a simple explanation for the fact that speakers are responsible for assertoric contents.

This approach parallels "fittingness" approaches to value and responsibility familiar from metaethics. The fittingness approach to value holds that to be valuable is to be a fitting object of favoring attitudes. The fittingness approach to responsibility analyzes one's being responsible for φ in terms of its being fitting to hold one responsible for φ.[13] The basic thought behind these approaches is that value and responsibility are both epistemologically and metaphysically mysterious. However, our practices of valuing things or holding people responsible are not especially mysterious. For example, it is clear that we do value things, and that we can be more or less justified in doing so. Thus, we can render value less mysterious by making sense of it in terms of our practice of valuing.

Assertoric content is similarly mysterious: It is hard to give a plausible theory of the grounds of assertoric content or the determination of contextual value. However, our practices of holding one another to particular claims, or particular word uses are not at all mysterious. So, by making sense of assertoric content in terms of these practices, we are able to demystify our metasemantics. Moreover, if the fittingness approach to responsibility is correct, then the word uses we are responsible for will be precisely those it is most fitting to hold us to. So, there will

---

[13]See Zimmerman (2010) for a discussion of the parallels between fittingness approaches to value and responsibility.

be a level of content we are responsible for that is directly determined by the word uses it is most fitting to hold us to on an occasion of use. This level of content is an obvious candidate for assertoric content.

Of course, this still leaves many questions unanswered. First, it is natural to ask, what is it to hold a speaker to a particular word use? Well, holding a speaker to a particular word use (and a corresponding claim) involves adopting a complex of attitudes and dispositions toward the speaker based on both the connotative aspect of the word's use, and the content determined by its use in its linguistic context. These attitudes will include simple reactive attitudes (such as negative appraisal if the word carries a harmful connotation, or if the corresponding claim is false), together with more complex dispositions and normative expectations (such as the expectation that the speaker will retract their claim if it is false, and that they will clarify their word use if the corresponding claim or connotation fails to correspond to their intention). Providing a full account of the complex of attitudes and dispositions involved in holding a speaker to a particular claim would require significant empirical input, and is beyond the scope of this article. However, I assume that most readers will be experienced enough communicators to have a good sense of what is involved.

Second, what is fittingness? If our hope is to demystify our metasemantics, then we had better be able to make good sense of this core notion. I, following the dominant trend in metaethics, am understanding fittingness in terms of reasons (see, e.g., Scanlon, 1998): The word use it is fitting to hold a speaker to on a given occasion is whatever word use there is the most objective reason to hold them to. There are many reasons for which we might hold a speaker to a particular use of a term, or to a particular claim. For example, the fact that a generic, expected, or actual audience member would interpret a speaker in a particular way will generally constitute a reason to hold a speaker to that interpretation. Likewise, the fact that the speaker intends to use a term in a particular way will constitute a reason to hold them to that use. These, and potentially innumerable other considerations will have to be weighed against one another in a given case in order to determine the use to which it is most fitting to hold a speaker. In some contexts, some such factors may be weighed more heavily than in others.[14]

---

[14]The reasons-based approach is the dominant approach to fittingness in metaethics (see, e.g., Scanlon, 1998). However, it is worth mentioning that this approach is subject to what is called the "wrong kind of reasons" problem. The problem is as follows: imagine some demon promised to destroy the world if we do not hold a speaker to a particular use of "that." This is a good reason to hold them to the relevant use of "that." But it does not make it fitting to hold the speaker to the relevant use of "that." It is the wrong kind of reason. There is a vast literature on the wrong kind of reasons problem in metaethics. Thus, there is not sufficient space to outline a full response here. However, since I imagine most readers will be unfamiliar with this literature I will mention the response to which I am most sympathetic. I believe that fittingness is only determined only by the *right* kind of reasons, and the right kind of reasons are object-based reasons:

> State based versus object-based reasons: The wrong kinds of reasons for an attitude are those that derive simply from one's holding the attitude in question. These are state based reasons. The right kinds of reasons derive from properties of the object of the attitude (in this case, the relevant utterance). (see Parfit, 2001; Piller, 2006).

So, the demon's promise provides a reason to hold the speaker to the claim in question. However, this reason derives not from features of the utterance itself, but rather from the consequences of holding the attitude (independently of any features of the utterance itself). There are many alternative responses to the wrong kind of reason problem which are consistent with the reasons-based approach to fittingness (Gibbard, 1990; Hieronymi, 2005; Lang, 2008; Rowland, 2013; Samuelsson, 2013; Schroeder, 2010; Skorupski, 2007; Stratton-Lake, 2005; Way, 2012).
It is worth noting that the wrong kind of reasons problem can also be avoided by treating fittingness as prior to reasons(see McHugh & Way, 2016; Zimmerman, 2010). I believe the fittingness view of contextual value works just as

We can see how such factors might be weighed to determine contextual value by returning to some of the cases we considered earlier:

WHALES: Smith uttered the sentence "I saw a sperm whale." When he used the term "sperm whale" he really meant "humpback whale." This counts in favour of holding him to a use of "sperm whale" according to which "sperm whale" picks out humpback whales. However, this reason is clearly outweighed by the fact that any normal audience member of the type Smith might encounter in this situation would take him to mean that he saw a sperm whale. In light of this, they would likely form the false belief that he saw a sperm whale. The blame for their false belief would land squarely on Smith. Thus, it is more fitting to hold Smith to the claim that he saw a sperm whale.[15]

STEAK: Martha uttered the sentence "my steak is raw." Raw, in English, means uncooked. This counts in favour of holding her to a use of "raw" according to which it means "uncooked." However, no reasonable individual in this situation would take Martha to be using "raw" literally. Moreover, she did not intend to be using "raw" literally, and she had good reason to expect that her audience was reasonable and competent. Thus, the fact that "raw" literally means uncooked does not count for much in this context.

X-RAY: Here the speaker knows a great deal about her audience. She knows that her audience has specialist knowledge not available to the average language user. Thus, the fact that an average language user would not understand her holds little weight when we consider what claim she should be held to. It is thus appropriate to hold her to her intended use of "that."

SPEAKS'S LECTERN: Here the speaker does not have an audience, thus the reaction of the actual audience holds no weight whatsoever. However, the speaker has many reasonable beliefs about his audience, and these reasonable beliefs affect what it is reasonable to hold him to. Moreover, he has a clear communicative intention which aligns with these reasonable expectations. Thus, the overall balance of reasons counts in favour of holding him to be referring to the lectern.

---

well on such "fittingness first" approaches. However, a number of my claims would have to be rephrased in terms of "fit making facts" rather than reasons.

[15]It might seem that the explanation I am giving here is circular: I have argued that we need a metasemantics that makes sense of our being committed to the content determined by the contextual values of the words used. But here it looks like I am suggesting that the contextual value of the words used is determined by the assertoric commitments that would be generated by this contextual value. But this is not quite right: There being most objective reason to hold the speaker to using "sperm whales" to pick out sperm whales *is* explained in part by the fact that any normal audience member would interpret him as intending to communicate that he saw a sperm whale. They would, thus, form a false belief. And this is something for which the speaker should be held responsible. Importantly, the notion of *assertoric* content is never evoked here—assertoric content is not simply the content that a normal competent audience member would take to be intended (this would be a version of the idealized interpreter view). The contents evoked in this explanation are *intended* content and, building on this, what a normal interpreter would take to be the intended content. So, the appearance of circularity is illusory.

> FRENCH INVASION: Fernando also had reasonable beliefs about his audience. The fact that his expected audience would interpret him as referring to 1066 is a good reason to hold him to this interpretation. However, his beliefs were badly mistaken. As a result (and through no fault of their own) his audience formed a false belief, and panic ensued. This, I believe, outweighs Fernando's reasonable beliefs: He was the one who put himself out there and took a risk with his audience's doxastic states. He reasonably expected this risk to be negligible. But, since he was wrong, and since his audience had no active role in these events, the buck ultimately stops with him.

It is also worth briefly considering cases of speaker error in which the audience is able to accommodate the error, as these cases might be thought to be problematic for the fittingness approach. Consider the following:

> GRANDMOTHER: A grandmother regularly mixes up the names of her dog Moomin, and her grandson Peter. However, it is almost always clear from the context who she intends to refer to. One day, whilst speaking to someone who knows her well, she says "Peter was burying a bone in the garden earlier." It is clear to the audience that she means to refer to her dog rather than her grandson. It might seem that the fittingness account makes the wrong prediction here. Intuitively, the assertoric content of her utterance is that her dog was burying a bone. Yet, there seems to be a sense in which it would be fitting to hold her to the claim that her grandson was burying a bone—a request for clarification would not be out of order. So, does the fittingness account make the wrong prediction here? No. According to the fittingness approach the contextual value a term receives on an occasion of use is the value it is *most* fitting to hold the speaker to. And this is the use that corresponds to the intuitive assertoric content of her utterance. It is *most* fitting to take her use of "Peter" to refer to her dog, and to hold her to this use. After all, a request for clarification would likely be met with "you know exactly who I meant," and it would be inappropriate to direct any negative reactive attitudes toward her on the basis of her grandson's not having buried a bone in the garden.

Some cases will be clearer than others. For example, I take X-RAY and SPEAKS'S LECTERN to be clearer than FRENCH INVASION. This is true with respect to intuitions about contextual value and assertoric content, along with our responsibility judgments. That is, in FRENCH INVASION it is less intuitively obvious that the Fernando's use of "now" refers to the present rather than 1,066, and it is also less clear which use he should be held to. Intuitions about such cases will, I believe, largely depend on the normative weight we are inclined to give to different competing considerations (e.g., how heavily should we weigh the fact that Fernando reasonably believed his audience to be in on the pretence?).

4 | CONCLUSION

I started with a puzzle: Once distinguished from intended content and standing content it is no longer clear why assertoric content is important. It seems to lack an obvious role. Yet it must play some important role in our communicative practices, otherwise we would not track

assertoric content, and it would not seem so central to our practices of appraising speech. In response I suggested that assertoric content provides a means for us to track and talk about the commitments we undertake when we speak. However, there are different levels of content we become responsible for through speech. For example, there is a sense in which we are responsible for the contents of our implicatures as well as our assertions. So, we need a way to distinguish assertoric commitments from other commitments. I suggested that we do so in terms of the manner in which these commitments are generated: Assertoric commitments are generated directly by the contextual values of the words used.

This raised two questions: (a) Why is the content determined in this way something that speakers make themselves responsible for? (b) Why is it important for us to distinguish between communicative commitments in terms of the manner in which they are generated? I have spent most of the paper answering this first question, and drawing out the lessons for metasemantics: I suggested that (a) assertoric content is the content it is fitting (i.e., there is most reason) to hold a speaker to in light of their utterance and (b) contextual value corresponds to the word use it is fitting (i.e., there is most reason) to hold a speaker to. This approach was shown to do better than its closest rival: the idealized interpreter approach. It was also shown to explain the appeal of both the idealized interpreter approach, and intentionalism. Indeed, not only does the fittingness view explain the appeal of its nearest rivals, it also captures the grain of truth in a thought often appealed to by those skeptical of traditional approaches to semantics (such as Pietroski, 2003, cited in Section 2.1): The truth conditions of a given utterance will be sensitive to innumerable factors that are resistant to systematic theorizing. As we have seen, the factors that determine the claim it is most fitting to hold a speaker to (and, thus the assertoric content of their utterance) will vary from context to context. So, in this sense, a fully systematic metasemantics does seem beyond the reach of science. However, this does not prevent us from understanding the determination of assertoric content, and explaining why particular terms receive particular values in particular contexts.

Having answered this first question we should, in closing, consider the second. Here is the question: There is a distinction between assertoric and other forms of content. This distinction seems important, we have a special set of concepts for tracking this distinction (e.g., the concepts of assertoric content, of implied content, of presupposed content, etc.). Moreover, assertoric content seems special in the following way: Our practices of appraising communicative acts seem to center on assertoric content. Many theorists assume that assertion, unlike other forms of communication (such as implicature), are committal; or, they at least assume that assertion is committal in a special normatively significant way. If they are right, then it would be unmysterious why we consider the distinction between assertoric and other forms of content to be so important. However, I (along with others: Hawthorne, 2012; Peet, 2015; Viebahn, 2017, 2020, forthcoming) have suggested that assertoric content does not carry with it a distinct form of commitment lacking from other speech acts. Rather, I have argued that the difference concerns the manner in which the commitments are generated. But why would it matter *how* the commitments are generated? If this is where the distinction between assertoric and other forms of content resides, why is this distinction important?

My answer is as follows: It is clearly important for us to be able to track the commitments speakers undertake when they communicate. However, this is really more important for audiences than it is for speakers. As speakers we will often (for good and bad reasons) wish to avoid being held to the things we have communicated. Moreover, even as audiences it sometimes suits us for there to be means of communication where the responsibilities undertaken are not so easily tracked: There will sometimes be information that speakers would be less likely to

share if they could easily be pinned down as having done so. Audiences would be shut off from this information if speakers did not have a less committal means of communication at their disposal. These competing demands pull us in different directions: There is pressure for us to develop a set of conceptual resources for tracking the commitments speakers undertake when they communicate. But there is also pressure for these resources to be imperfect.

One way of achieving this is by drawing a distinction between different ways of generating communicative commitments, and centering our practices of appraising speech on one such manner of generating communicative commitments. This is reflected in the set of practices and conceptual resources we have available for appraising communicative acts. It is true that we have concepts such as "insinuation" and "presupposition" for tracking non-assertorically communicated contents. But far more central to our practices of reporting speech, and far more easily available to the average language user, are the resources we use to track assertoric commitment. Moreover, it is these commitments to which our communicative norms, and our most damning appraisals of defective speech (e.g., "lie") most obviously apply. There is, in this sense, an imbalance in the resources we have available for tracking and holding speakers to the commitments they undertake when they speak. This imbalance places assertoric commitments above all others. And it reflects the conflicting pressures governing the development of the resources we have for tracking communicative commitments.

But surely, there are other ways in which our resources for tracking communicative commitments could have developed. Why do we distinguish between communicative commitments specifically in terms of their manner of generation? Well, the choice could be arbitrary. There is pressure for our commitment tracking resources to develop in such a way that there is some such imbalance. Perhaps, it is just by chance that our practices developed as they did. But think more can be said in favor of distinguishing between communicative commitments in terms of their manner of generation: Doing so gives us a straightforward method for generating easily trackable and not so easily trackable commitments: to generate easily trackable commitments, generate your commitments directly. To generate difficult to track commitments, generate your commitments indirectly.

DATA AVAILABILITY STATEMENT

Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

ORCID

Andrew Peet https://orcid.org/0000-0001-9143-2201

REFERENCES

Armstrong, J. (2016). The problem of lexical innovation. *Linguistics and Philosophy*, *39*, 87–118.

Bach, K. (1994). Semantic slack: What is said and more. In S. Tsohatzidis (Ed.), *Foundations of speech act theory: Philosophical and linguistic perspectives* (pp. 267–291). London: Routledge.

Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.

Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.

Borg, E. (2019). Explanatory roles for minimal content. *Noûs*, *53*(3), 513–539.

Brandom, R. (1983). Assertion. *Noûs*, *17*(4), 637–650.

Brandom, R. (1994). *Making it explicit*. Cambridge, MA: Harvard University Press.

Camp, E. (2018). Insinuation, common ground, and the conversational record. In D. Fogal, D. Harris & M. Moss (Eds.), *New work on speech acts* (pp. 40–67). Oxford: Oxford University Press.

Cappelen, H. (2011). Against assertion. In J. Brown & H. Cappelen (Eds.), *Assertion: New philosophical essays* (pp. 21–49). Oxford: Oxford University Press.

Cappelen, H. (2019). Assertion: A defective theoretical catagory. In S. Goldberg (Ed.), *Oxford handbook of assertion* (pp. 139–159). Oxford: Oxford University Press.

Cappelen, H. & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.

Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.

Chomsky, N. (1995). Language and nature. *Mind*, *104*(413), 1–61.

Dummett, M. (1973). *Frege: Philosophy of language*. London: Gerald Duckworth.

Evans, G. (1979). Reference and contingency. *The Monist*, *62*(2), 161–189.

Fricker, E. (2012). Stating and insinuating. *Proceedings of the Aristotelian Society*, *86*, 61–94.

Gibbard, A. (1990). *Wise choices, apt feelings*. Oxford: A theory of normative judgment, Clarendon Paperbacks.

Grice, P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Hawthorne, J. (2012). Some comments on Fricker's stating and insinuating. *Proceedings of the Aristotelian Society*, *86*(1), 95–108.

Heck, R. K. (2014). Semantics and context dependence: Towards a Strawsonian account. In B. Sherman & A. Burgess (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 327–364). Oxford: Oxford University Press.

Hieronymi, P. (2005). The wrong kind of reasons. *Journal of Philosophy*, *102*(9), 437–457.

King, J. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, *27*, 288–311.

King, J. (2014a). Speaker intentions in context. *Noûs*, *48*(2), 219–237.

King, J. (2014b). The metasemantics of contextual sensitivity. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 97–118). New York, NY: Oxford University Press.

Lang, G. (2008). The right kind of solution to the wrong kind of reason problem. *Utilitas*, *20*(4), 472–489.

Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Ohman (Eds.), *Philosophy and grammar* (pp. 79–100). Amsterdam: Reidel.

Ludlow, P. (2014). *Living words: Meaning underdetermination and the dynamic lexicon*. Oxford: Oxford University Press.

MacFarlane, J. (2011). What is assertion? In H. Cappelen & J. Brown (Eds.), *Assertion: New philosophical essays* (pp. 79–96). Oxford: Oxford University Press.

McHugh, C. & Way, J. (2016). Fittingness first. *Ethics*, *126*(3), 575–606.

Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, *33*(5), 335–380.

Nowak, E. & Michaelson, E. (2021). Who's your ideal listener. *Australasian Journal of Philosophy*, *99*(2), 257–270.

Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, *3*, 143–184.

Parfit, D. (2001). Rationality and reasons. In D. Egonsson, J. Josefsson, B. Petersson & T. Rønnow-Rasmussen (Eds.), *Exploring practical philosophy: From action to values* (pp. 17–39). Aldershot: Ashgate.

Peet, A. (2015). Testimony, pragmatics, and plausible deniability. *Episteme*, *12*(1), 29–51.

Peet, A. (2019). Knowledge-yielding communication. *Philosophical Studies*, *176*, 3303–3327.

Perry, J. (2006). Using indexicals. In M. Devitt & R. Hanley (Eds.), *The Blackwell guide to the philosophy of language* (pp. 314–334). Oxford: Blackwell.

Pietroski, P. (2003). The character of natural language semantics. In A. Barber (Ed.), *The epistemology of language* (pp. 217–256). Oxford: Oxford University Press.

Piller, C. (2006). Content-related and attitude-related reasons for preferences. *Royal Institute of Philosophy Supplements*, *81*(59), 155–181.

Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, *189*(1), 75–96.

Rabern, B. (2017). A bridge from semantic content to value. *Philosophical Topics*, *45*(2), 201–226.

Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.

Romdenh-Romluc, K. (2002). Now the French are invading England! *Analysis*, *62*(1), 34–41.

Romdenh-Romluc, K. (2006). I. *Philosophical Studies*, *128*(2), 257–283.

Rowland, R. (2013). Wrong kind of reasons and consequences. *Utilitas*, *25*(3), 405–416.

Samuelsson, L. (2013). The right version of 'the right kind of solution to the wrong kind of reason' problem. *Utilitas*, *25*(3), 383–404.

Scanlon, T. (1998). *What we owe each other*. Cambridge, MA: Harvard University Press.

Schroeder, M. (2010). Value and the right kind of reason. *Oxford Studies in Metaethics*, *5*, 25–55.

Shapiro, L. (2020). Commitment accounts of assertion. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 73–97). Oxford: Oxford University Press.

Skorupski, J. (2007). Buck-passing about goodness. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, & D. Egonsson (Eds.), *Hommage à Wlodek*. Philosophical Papers Dedicated to Wlodek Rabinowicz.

Speaks, J. (2016). The role of speaker and hearer in the character of demonstratives. *Mind*, *125*(498), 301–339.

Sperber, D. & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Blackwell.

Stainton, R. (2016). Full-on stating. *Mind & Language*, *31*(4), 395–413.

Stanley, J. (2002). Modality and what is said. *Philosophical Perspectives*, *16*, 321–344.

Stonjić, U. (2017). On the connection between semantic content and the objects of assertion. *Philosophical Topics*, *45*(2), 163–178.

Stratton-Lake, P. (2005). How to deal with evil demons: Comment on Rabinowicz and Rønnow-Rasmussen. *Ethics*, *115*(4), 788–798.

Viebahn, E. (2017). Non-literal lies. *Erkenntnis*, *82*, 1367–1380.

Viebahn, E. (2020). Lying with presuppositions. *Noûs*, *54*(3), 731–751.

Viebahn, E. (forthcoming). The lying–misleading distinction: A commitment based approach. *The Journal of Philosophy*.

Way, J. (2012). Transmission and the wrong kind of reason. *Ethics*, *122*(3), 489–515.

Wettstein, H. (1984). How to bridge the gap between meaning and reference. *Synthese*, *58*, 63–84.

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 17–54). Oxford: Oxford University Press.

Zimmerman, M. (2010). Responsibility, reaction, and value. *The Journal of Ethics*, *14*(2), 103–115.