

Åsa Wikforss¹
Stockholm University, Sweden

ARE THERE UNDERSTANDING-ASSENT LINKS?

Abstract: It is commonly held that there are internal links between understanding and assent such that being semantically competent with an expression requires accepting certain sentences as true. The paper discusses a recent challenge to this conception of semantic competence, posed by Timothy Williamson (2007). According to Williamson there are no understanding-assent links of the suggested sort, no internal connection between semantic competence and belief. I suggest that Williamson is quite right to question the claim that being semantically competent with an expression *e* requires accepting a certain sentence *S* as true. However, Williamson does not merely wish to reject this version of the understanding-assent view, but the very idea that the connection with belief provides constitutive constraints on linguistic understanding and concept possession. This further move, I argue, is very problematic. Giving a plausible account of semantic competence requires accepting that there are constitutive links between understanding and assent, although these links should be construed holistically rather than atomistically.

1. INTRODUCTION

It is much disputed what I know when I know a language. Is knowledge of language a form of 'knowledge that', as cognitivists would have it, or does it reduce to knowing how? A common objection to cognitivism is that it, implausibly, requires the speaker to have a rich set of beliefs about the semantics (and syntax) of her language, i.e. a set of meta-beliefs. If knowledge of language is a mere ability, by contrast, no such requirement is involved – the speaker could, as it were, be 'ignorant' of her own language.¹ Proponents of both camps, however, tend to agree that there is another sense in which knowledge of language, linguistic understanding, involves belief: Being semantically competent with an expression involves accepting certain sentences as true, it involves having certain (object-level) beliefs. This, it is held, is *constitutive* of semantic competence, of understanding. Hence, if *P* does not accept

¹ The paper is based on comments given to Tim Williamson's Wedberg lectures in Stockholm, April 2006, and benefited much from Williamson's very helpful response. Thanks also to Kathrin Glüer, Sören Häggqvist and Peter Pagin for valuable comments on an earlier version of the paper.

sentence *S* as true, she manifests lack of understanding of some of the component expressions.ⁱⁱ

In the paper I discuss a recent challenge to this conception of semantic competence, posed by Timothy Williamson (2007). Williamson characterizes the target theses as follows, one for the case of language, the other for thought:

(UAI) Necessarily, whoever understands the sentence 'Every vixen is a female fox' assents to it (73).

(UAt) Necessarily, whoever grasps the thought *every vixen is a female fox* assents to it (74).

Both theses, according to Williamson, are false: There are no understanding-assent links of the suggested sort, no constitutive links between semantic competence and belief. The idea that there are such links, Williamson argues, rests on a mistaken conception of semantic competence. In its place we should put an alternative conception, one that cuts the link with belief and is based on externalist considerations.

I suggest that Williamson is quite right to question theses (UAI) and (UAt). However, it is clear from Williamson's discussion that he does not merely want to reject this version of the idea that there are internal links between understanding and assent (the 'UA-thesis', for short). He also denies that there are understanding-assent links of *any* sort, and suggests that we should reject the very idea that the connection with belief provides constitutive constraints on linguistic understanding and concept possession. This, I argue, is much more problematic. Williamson's own alternative account, falling back on externalist considerations, cannot give a satisfactory account of semantic competence. To give such an account, I suggest, we need to recognize that there are constitutive links between understanding and assent, although these links should be construed holistically, and not along the lines of (UAI) and (UAt).

First, some preliminary remarks. When discussing knowledge of meaning it is helpful to start from a fairly neutral, 'minimalist' conception of such knowledge. On this conception, P grasps the meaning m of expression e if it is true that e means m for P . Construed this way, the question of semantic knowledge is not distinct from the foundational question concerning the *determination* of meaning: In virtue of what facts does e mean m for P , and how is the function from these facts to meanings to be understood? Whatever theory is adopted concerning the relevant meaning determining facts (whether one opts for facts about P 's use of e , facts about P 's causal interaction with the environment, facts about conventions, etc.), if the facts that determine that e means m for P are in place, then P knows the meaning of e . It is then an open question whether this knowledge should be construed as a form of propositional knowledge or as a practical ability.

Although this minimalist conception of knowledge of meaning is neutral with respect to the debate over cognitivism, it should be noted that it is not neutral in another respect: It presupposes that the conditions for knowing the meaning of e , understanding that e means m , coincides with the conditions for meaning m by e . According to some theories of understanding, this is not the case. In particular, as we shall see below, some versions of externalism imply that it may be true that e means m for P , although P fails to (fully) understand the meaning of e .

2. THE UNDERSTANDING-ASSENT VIEW

The claim that there are internal links between understanding and assent is commonly made against the backdrop of Quine's criticisms of the analytic-synthetic distinction. It is granted that Quine has shown that some conceptions of analyticity are problematic, in particular the idea that there are truths 'in virtue of meaning alone', what is sometimes called metaphysical analyticity. However, it is argued, if we are to account for semantic competence we still need

to appeal to a form of analyticity, so called *epistemological analyticity*.ⁱⁱⁱ On this view, if *P* understands *e*, she must assent to a set of privileged sentences – those that are constitutive of competence with expression *e*. A sentence belonging to this privileged set is analytic in the epistemological sense. On the further assumption that knowledge of meaning is a priori, it is argued, the speaker can know *S* a priori, without having to rely on experiential evidence.^{iv}

The notion of assent, in this context, is that of a mental attitude, i.e. belief. To say that there are internal links between understanding and assent, therefore, is to say that there is a constitutive link between understanding, grasp of meaning, and belief. The appeal to understanding-assent links therefore limits the scope of disagreement in belief – a disagreement on the meaning constitutive sentences must be construed as a disagreement in meaning (concepts). Given the further connection between belief and sincere assertion, the link between understanding and assent transfers to a link between understanding and use.^v For example, a speaker who sincerely asserts 'Not all vixens are female foxes' does not understand one of the component expressions.^{vi}

The idea that in order to account for semantic competence it is required that we appeal to a notion of analyticity can be found already in Grice and Strawson's famous response to Quine.^{vii} They give as example a speaker, *Y*, who asserts the following two sentences:

- (i) 'My neighbor's three-year-old child understands Russell's Theory of Types'
- (ii) 'My neighbor's three-year-old child is an adult'

We would respond to (i) with disbelief, they argue, but if the child were produced and expounded the theory correctly etc. we would accept it as true, even if amazing. However, our reaction to (ii) would be very different. If *Y* insists that he is not speaking figuratively, but is making a sincere assertion then the proper conclusion is that he just does not know the meaning of some of the words used. According to Grice and Strawson, this illustrates the

availability of an informal explanation of the notion of analyticity. If someone rejects a highly plausible empirical truth we will simply disagree with him, whereas if someone rejects an analytic truth we will not *understand* him, and will have to conclude that he does not understand what he is saying.

Williamson's criticism of (UAI) and (UAt) is also connected with the analyticity debate. His purpose is meta-philosophical and he sets out to argue against the claim that philosophical truths are analytic or conceptual truths. His rejection of understanding-assent links is therefore primarily directed against epistemological conceptions of analyticity.

However, one may be a UA-theorist without committing to epistemological analyticity. To see this, it is important to keep clear on the distinction between semantics and meta-semantics, or foundational semantics. Sometimes theses such as (UAI) and (UAt) are put forth within the context of *meta-semantics*, telling us something about the *determination* of meaning and content. On this view, what determines the meaning of 'vixen' is the fact that the speaker is disposed to assent to a certain set of sentences where the expression is a component, such as 'Every vixen is a female fox'. This does not imply that 'Every vixen is a female fox' expresses an analytic truth since it does not imply that the sentence is constitutive of the *meaning* of 'vixen', just that it is constitutive of *grasp* of this meaning. Thus, it is perfectly possible to hold that there are understanding-assent links in the case of sentences that are not plausible candidates for analytic truths but straightforwardly empirical and a posteriori.^{viii} Moreover, as long as the UA-thesis is presented as a meta-semantic thesis, it does not even follow that the sentence assented to is *true*. It could be claimed that it is constitutive of understanding that one assents to *S*, even while *S* is false. For instance, it has been argued by Matti Eklund that our competence with certain terms (such as 'true') requires us to be disposed to accept certain sentences as true that are not in fact true.^{ix} Another example might be provided by the debate over vagueness, where it is commonly claimed that

it is part of one's semantic competence with vague words to accept the principle of tolerance – a principle everybody agrees must be false (on pain of paradox). Also, it is quite possible to include among the facts that serve to determine meaning the disposition to accept certain empirical falsehoods in situations where error reliably occurs.^x

However, (UAI) and (UAt) could also be employed to provide a *semantic* theory, telling us something about the *meanings* of the relevant expressions, rather than about the determination of meaning. For instance, it may be held that the meaning of certain terms (such as the logical constants and theoretical terms) is constituted by their inferential role.^{xi} If so, there will be a direct connection with analyticity since it follows that certain sentences will be constitutive of the meaning of *e*: these sentences, then, will be true and they will (arguably) be such that they can be known a priori, simply on the basis of understanding the relevant expression.^{xii}

Although Williamson's main target is epistemological analyticity, it is clear that he wishes to reject both the semantic and the meta-semantic version of the UA-thesis. He rejects inferential role semantics but he also agrees on the meta-semantic level, offering a version of externalism that cuts the links between meaning and belief. In what follows, I shall focus on the meta-semantic question, and on Williamson's claim that we need not assume the existence of any type of understanding-assent links in our meta-semantics.

3. WILLIAMSON AGAINST UNDERSTANDING-ASSENT LINKS

Williamson argues against (UAI) and (UAt) by way of example. He examines cases of suggested understanding-assent links, and argues that the links may fail and yet the speaker should be described as a perfectly competent speaker. Thus, in response to Grice and Strawson's example, he suggests that a competent speaker may well believe that some three-year olds are adults, 'explaining away all the evidence to the contrary by *ad hoc* hypotheses or

conspiracy theories (many three-year-olds pretend to be eighteen-year-olds in order to vote, the abnormally polluted local water slows development, and so on) (2007: 85)'. Although we may consider the person foolish and obviously wrong, Williamson argues, it does not follow that she is semantically incompetent.

Worrying that someone may object to this particular example, Williamson imagines an even simpler case: an individual, Peter, who doubts a very simple logical truth. Peter has odd views and believes that there is a logical entailment from 'Every F is a G' to 'There is at least one F'. Since, moreover, he has been spending too much time on the Internet he has come to believe that there are no vixens. Consequently, he doubts the truth of 'All vixens are vixens' (86). However, Williamson argues, Peter is a perfectly competent speaker. Hence, the alleged links between understanding and assent do not even hold in the case of simple logical truths.^{xiii}

Williamson's argument, notice, does not draw on the distinction between competence and performance.^{xiv} The claim is not that Peter's competence is in some sense blocked, that he is really disposed to assent to *S*, and makes a performance error. Rather, the claim is that Peter has no disposition whatsoever to assent to *S*, as a result of his theoretical commitments (99, ff.). Peter's failure to assent to *S* does not in any way reflect on his linguistic competence since, according to Williamson, there simply are no such links between competence and assent in the first place.

Moreover, Williamson's claim is not that deviant speakers display an *incomplete* understanding of the meaning of *e*. In this respect Williamson departs from social externalists, such as Tyler Burge, who argue that deviant speakers should be ascribed an incomplete understanding of the meaning of the expression in question. Thus, Burge stresses that when the speaker (in the actual world) asserts 'I have arthritis in my thigh' he makes a conceptual error and displays an incomplete grasp of the meaning of 'arthritis'. Still, Burge argues, in

such a case the meaning of the word expressed should be interpreted in accordance with the community practice. This means that although Burge (like Williamson) questions the link between *meaning determination* and assent, he accepts the thesis that there are internal links between *understanding* and assent. Burge does therefore not subscribe to the minimalist conception of knowledge of language, mentioned above: It may be true that 'arthritis' means *arthritis* for *P*, even though *P* does not (fully) understand the meaning of 'arthritis'.

Williamson, by contrast, rejects the appeal to incomplete understanding. The speakers in his examples, he stresses, are *experts* – their understanding of the relevant meanings is as complete as it can be: "experts themselves can make deviant applications of words as a result of theoretical errors and still count as fully understanding their words" (98). Moreover, he suggests, even if they were not experts, there would still not be any reason to treat them any differently and ascribe incomplete understanding to them (99). Unlike Burge, therefore, Williamson sticks to the minimalist conception of semantic competence.

Instead of appealing to the idea of blocked competence dispositions, or the notion of incomplete understanding, Williamson appeals to Quine's epistemological holism. The epistemological status of one's beliefs, Williamson suggests, depends on its position in a larger network of beliefs. This implies that there can be no 'litmus test' for understanding, since someone could fail it and yet count as being semantically competent in virtue of her overall use with the term. Thus, deviances at one point can be compensated for by conservatism on others: 'Epistemological holism explains how unorthodoxy at one point can be compensated for by orthodoxy on many others' (2007:91).

Williamson's challenge to (UAI) and (UAT) should, I believe, be taken seriously. There is a strong intuition, in the cases described, that the speaker in question does not manifest any kind of linguistic incompetence. The speakers are generally competent, and their disagreement is limited in the sense that it does not spread to large bodies of related

belief and does not prevent communication.^{xv} For instance, the speaker who claims that the neighbor's three-year old is an adult could still use 'adult' in perfectly standard ways otherwise (holding that in the overwhelming majority of cases three-year olds are not adults, etc.). It is simply implausible to insist that such isolated disagreements *must* involve semantic incompetence. This much we have learned from Quine.^{xvi} Moreover, to fall back on the notion of 'blocked competence' seems perfectly ad hoc, since there is no evidence that the speaker has the relevant disposition in the first place, for instance, that Peter has the disposition to accept 'Every vixen is a female fox'.

Similarly, I think Williamson is quite right to reject the appeal to incomplete understanding. As noted above, the idea that speakers have incomplete understanding of the meaning of their words (the concepts expressed) requires that we reject the minimalist conception of semantic competence, separating the conditions of meaning *m* by *e* from the conditions of understanding *e*. To sustain such a separation one would have to explain what is required for full understanding, beyond simply meaning *m* by *e*, and it is not clear how this is to be done.^{xvii} Moreover, even if one accepts that there is some such distinction, and that there are cases of incomplete understanding, the question remains whether this is plausible in the cases under consideration. Just as we should be skeptical of the suggestion that isolated disagreements entail semantic incompetence, we should be skeptical of the suggestion that speakers who disagree with the community on single points (for instance, on the truth of the sentence 'Arthritis afflicts the joints only') should be described as having incomplete understanding of meaning (or, for that matter, as using the word with a non-standard meaning).^{xviii} Every such deviance can be made sense of against the larger background of the speaker's beliefs and need not entail lack of understanding.

The question is where this leaves us. Assuming that the simple understanding-assent view fails, what should we put in its place? After all, as Williamson notes, it cannot be

denied that there is a distinction between understanding a word and not understanding it (126). Assuming the minimal view of semantic competence, again, this is just to say that we should be able to distinguish the case where *P* means *m* by *e* from the case where *P* does not mean *m* by *e*. Of course, there may be indeterminacies, and even when there is not we may be a long way from formulating a complete meta-semantic theory. However, if Williamson's rejection of the meta-semantic UA-thesis is to carry conviction, we need to be given some reasons to believe that semantic competence can be accounted for without appealing to understanding-assent links.

4. THE ALTERNATIVE?

One option would be to provide a modified (UA)-thesis, one that does not commit us to (UA1) and (UA2). Williamson, however, does not consider this option.^{xix} Instead, he appeals to a form of externalism, one that does not depend on the idea that there is a constitutive connection of any sort between meaning and belief. Semantic externalism, he suggests, explains how deviant speakers 'can still use the terms with their normal public sense' (91). What is required is merely 'enough connection in use' between the speakers – if this goes missing, there is no understanding. Failure of understanding, thus, consists in a failure to causally interact with the social practice: "One can lack understanding of a word through lack of causal interaction with the social practice of using that word, or through interaction too superficial to permit sufficiently fluent engagement in the practice" (126).

However, it is unclear how this appeal to the causal interaction with the community can provide an answer to the question of what it is to understand the meaning of a word, to be semantically competent. Of course, if *P* does not understand the community language, then the causal interaction with the community's linguistic practice will be rather limited. But causal interaction with the community simply does not seem necessary for

understanding the language.^{xx} Nor does it seem sufficient, unless, of course, the relevant interaction is that of *communication*, i.e. mutual understanding – but that would seem to get us no further if we seek an account of understanding in the first place.

Moreover, the appeal to causal interaction is of no help when it comes to illuminating *concept grasp*: Individuals may grasp the same concept without ever having causally interacted with one another. Williamson is fully aware of this. Sameness of concept, he stresses, does not entail causal relatedness (127). How, then, are we to distinguish the case where *P* does grasp the concept *vixen* from the case where *P* does not? Williamson, again, rejects the appeal to understanding-assent links and suggests a minimalist account of concept grasp. On the simplest view, he says, 'thinking a thought with any attitude towards it suffices for grasping it' (74). This minimalist claim parallels the minimalist conception of semantic competence, and it seems incontrovertible: If one thinks a thought then one grasps it. However, we still need an explication of what it is to grasp a concept in the first place, to think a thought with a certain content. And here, again, the appeal to the community is of no help.

It is therefore a serious question whether Williamson's appeal to externalist considerations provides a viable alternative theory of semantic competence and understanding, or even the outlines of such a theory.^{xxi} Indeed, it is difficult to see how the suggested picture could be a version of social externalism in the first place. According to standard social externalism meaning is determined by the experts' practice. This means that the experts themselves cannot make deviant applications of their words: a 'deviant' application would just be a case where the word has a different meaning.^{xxii} On Williamson's view, by contrast, experts can make deviant applications without it following that the word has a different meaning. What, then, determines meaning on Williamson's view? The idea, it seems, is that it is neither the individual's practice, nor the experts', but the practice of the community

as a whole. Each individual, including the experts, uses words as words of a public language: 'their meanings are constitutively determined not individually but socially, through the spectrum of linguistic activity across the community as a whole' (2007:98).

However, the appeal to 'the community as a whole' does not help. The relevant notion of a community is that of a *linguistic* community, i.e. a group of people speaking the same language. But what we wanted to know is, precisely, when a group of people could be said to speak the same language. Using the same words, obviously, is not sufficient since two speakers may use the same words with different meanings. That is, some deviances do constitute meaning disagreements and Williamson does not provide any means of determining which these are.

In fact, social externalism is normally presented as a form of UA-theory. As noted above, Burge questions the link between meaning determination and assent since he holds that there are cases of linguistic deviances where the individual should nonetheless be interpreted in accordance with the communal language (although, again, Burge does suggest that in those cases the individual should be ascribed an incomplete understanding of the expression in question). However, at the same time Burge is careful to make clear that this only holds for *some* cases and that there are deviances that are so radical that reinterpretation is called for. The notion of a 'radical divergence', on Burge's view, turns on the idea that there are rationality constraints on interpretation, limiting acceptable disagreement in belief. For instance, he suggests that when it comes to malapropisms and slips of the tongue, the principle of charity requires a non-standard interpretation: The individual who sincerely claims to have been drinking orangutans for breakfast is better interpreted as meaning *orange juice* by 'orangutan' than as using the word with its standard meaning.^{xxiii} This means that Burge's social externalism does, after all, rest on a version of the (UA)-thesis.

On Williamson's version of externalism, by contrast, there is no suggestion as to how we are to separate the case where a deviance requires reinterpretation from a case where it does not. Since there are *no* constitutive links between meaning and sincere assertions on his view, two speakers could use the word 'vixen' in radically different ways and yet both mean *vixen* by it. As a result, it is unclear what the relevant meaning determining facts are and how the function from these facts to meanings is to be understood. However, this means that Williamson has failed to provide us any reason to think that there is a viable account of semantic competence that does not depend on some version of the UA-thesis.

In fact, in his discussion of deviant speakers Williamson himself seems to rely on the existence of understanding-assent links. Consider how he presents the examples of these speakers: In each case we are told an elaborate story that explains the deviance. Thus, in response to Grice and Strawson's example, the person who assents to 'My neighbor's three-year-old child is an adult', Williamson argues that this need not display a failure to understand the component words since *P* may have some odd theory about pollution or some conspiracy theory. Similarly, in the case of Peter, we are told about Peter's 'foolish' beliefs about universal quantification. This suggests that there is, after all, a constitutive connection with belief, only a holistic one; it suggests that although we cannot single out single sentences as a litmus test for understanding, unless a plausible background story emerges such that the 'unorthodoxy is compensated for by orthodoxies at other points', the fact that *P* assents to *S* does entail failure of understanding. At some point, the strain on belief becomes too great, making sense of the speaker becomes too difficult, and it has to be concluded that *P* fails to understand some of the expressions involved.^{xxiv}

Williamson claims that the need for a background story is merely *practical*, it makes for ease of understanding but does not imply any internal link between understanding and belief. He grants that disagreement is more fruitful against a background of extensive

agreement, but denies that this imposes any constitutive constraints: "A practical constraint on useful communication should not be confused with a necessary condition for literal understanding" (125). But the suggestion that we are dealing with a merely practical constraint sits badly with Williamson's overall reasoning. If the connection with belief were merely practical and did not provide constitutive constraints of any sort, we should be able to dispose of the background story altogether and there would be no principled reason why the speaker could not be 'unorthodox throughout'.

At points, Williamson suggests that we could in fact dispose of the background story. For instance, he says that we can also imagine untheoretical native speakers whose patterns of assent and dissent are just like those of deviant logicians (2007: 99). Once we allow that the deviant logician is a competent speaker, Williamson suggests, we can hardly refuse the same classification of the untheoretical, deviant speaker as semantically competent. However, this carries little conviction. Imagine a speaker who walks around questioning sentences such as 'Every vixen is a female fox', and related inferences, without providing any reasons whatsoever for her deviance. Or a speaker who assents to 'The neighbor's three-year old is an adult' without appealing to any conspiracy theory or the like to justify her claim. Is there any inclination at all to say that such a speaker is semantically competent? I should think not, and the reason is precisely that there is not a set of background beliefs that allows us to make sense of the deviance.

The proper lesson of Williamson's criticisms of (UAI) and (UAt), I would therefore propose, is not that there are no understanding-assent links but that the links are holistic in nature. That the understanding-assent links may fail in any particular case does not imply that there are no such links. On the contrary, there are strong reasons to think that we cannot account for semantic competence without appealing to the idea that the connection with belief provides constitutive constraints on the determination of meaning and content. The

relevant holism, thus, is not merely epistemic holism, but *semantic* holism, understood as a theory within foundational semantics. Consider Burge's 'orangutan'-example. To mean *orangutan* by 'orangutan', on the holist view, *P* need not hold any particular belief about orangutans, but she does need to use the word 'orangutan' in such a way that interpreting her word as meaning *orangutan* makes her come out as having a reasonable set of beliefs about orangutans. The fact that the speaker assents to 'I drink orangutans for breakfast', therefore, does not automatically disqualify her as a speaker of English. She may have unusual habits (working at a zoo) or be sorely mistaken (a quack has sold her a very expensive miracle cure that he claims to be orangutan juice). The holism, that is, allows for (surprising) true beliefs as well as for errors – even errors such as those of Peter and Stephen. But when the would-be errors are radical, and no plausible background story emerges, it should be concluded that *P* does not understand the expression in the standard way.

There is a fear commonly associated with semantic holism: that it will undermine the very distinction that we wanted to secure, between meaning differences and belief differences.^{xxv} It is in order to prevent this that many have thought it necessary to appeal to a special set of meaning constitutive sentences. However, as long as the holism is meta-semantic the fear is unwarranted. What is required is merely that the function from use (assent) to meaning (concepts) is construed as a *many-one* function, such that several different patterns of use can determine the same meanings.^{xxvi} In addition, of course, we need to appeal to a principle that maps the meaning determining facts on to meanings; i .e. a principle that allows us to separate patterns of assent (and dissent) that determine the same meaning (for instance, that 'orangutan' means *orangutan*) from patterns that determine another meaning (that 'orangutan' means *orange juice*). Although I cannot discuss this further here, it is very likely that such a principle will have to involve rationality considerations, along the lines of the principle of charity, since this provides a tool for separating acceptable disagreements in

belief from unacceptable ones. As noted above, this is precisely the type of principle that Burge falls back on in order to account for the distinction between deviances that do require reinterpretation and those that do not.

The alternative to (UA1) and (UA2), I therefore propose, is not to cut the link between understanding and assent but to construe it holistically. This allows us to agree with Williamson that there is no litmus test of understanding, that in every particular case we can imagine a competent speaker who doubts the truth of *S*. However, the constitutive connection with belief ensures that unless a plausible set of background beliefs emerges, *P*'s use manifests lack of understanding.

References

- Burge, T. 1979. "Individualism and the Mental", *Midwest Studies in Philosophy* 4: 73-121.
—, 1986. "Intellectual Norms and the Foundations of Mind", *Journal of Philosophy* 83: 697-720.
—, 1993. "Concepts, Definitions, and meaning", *Metaphilosophy* 24: 309-325.
- Boghossian, P. 1996. 'Analyticity Reconsidered', *Nous* 30: 360-391.
—1997, 'Analyticity', in R. Hale and C. Wright eds., *A Companion to the Philosophy of Language*. Oxford: Blackwell.
—'Williamson on the *A Priori* and the Analytic', forthcoming in *Philosophy and Phenomenological Research*.
- Davidson, D. 1984. 'On Saying That', in *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
- Devitt, M. 2006. *Ignorance of Language*. Oxford: Oxford University Press.
- Eklund, M. 2002. 'Inconsistent Languages', *Philosophy and Phenomenological Research* 64: 251-75.
- Glüer, K. 2003. 'Analyticity and Implicit Definition', *Grazer Philosophische Studien* 66: 37-60.
- Grice, H.P. and Strawson, P.F. 1956. 'In Defence of a Dogma', *Philosophical Review* 65: 141-58.
- Pagin, P. 1997. 'Is Compositionality Compatible with Holism', *Mind and Language* 12.

Putnam, H. 1975. 'The Meaning of 'Meaning'', in *Minnesota Studies in the Philosophy of Science*, vol. 2, ed. K. Gunderson, Minneapolis: University of Minnesota Press: 215-271.

Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

Wittgenstein, W. 1969. *On Certainty*. Oxford: Blackwell.

References to the author taken out for blind refereeing.

ⁱ Devitt 2006.

ⁱⁱ The two camps can agree on this but will of course give a different diagnosis of the failure: On the cognitivist construal, the failure to assent to *S* reflects lack of a certain propositional knowledge, of certain meta-beliefs, on the non-cognitivist construal it reflects lack of the relevant knowing how (the ability to use *e* correctly, in accordance with its meaning).

ⁱⁱⁱ For the distinction between metaphysical and epistemological analyticity see Boghossian 1996 and 1997.

^{iv} See Boghossian 1996 and 1997.

^v The relation between assent construed as a mental attitude and assertion, of course, is complex. Even if the speaker only were to make sincere assertions, and would not make any performance errors, she may still lack a disposition to assert many things she believes. For instance, as Williamson suggests, she may find it embarrassing to assert trivialities such as 'Every vixen is a female fox' (74).

^{vi} Alternatively, assuming the minimalist conception of semantic competence, one of the component words has a non-standard meaning in *P*'s language.

^{vii} Grice and Strawson 1956: 150-152.

^{viii} For instance, there is the idea that if a speaker dissents from certain very basic and obvious empirical statements, such as 'This is a hand', she manifests lack of linguistic understanding. This is a leading theme of Wittgenstein's 'On Certainty' (see for example paragraphs 52-56).

^{ix} Eklund 2002: 253. However, Eklund overstates the case, moving from the meta-semantic claim to a semantic claim. Thus, he suggests that his account of competence shows that a principle (a sentence or an inference) can be constitutive of meaning, without being true (ibid: 256, 263). But this does not follow. Although it can be constitutive of *understanding* of meaning that one accepts a false sentence, a falsity cannot be constitutive of *meaning*.

^x This presupposes that the determination relation is a many-one function, i.e. that the relation between meaning determining facts and meanings is not an equivalence relation. I return to this issue below.

^{xi} The two may of course be combined: i.e. one may adopt an inferentialist semantics as well as an inferentialist meta-semantics. (This seems to be the view defended by Boghossian 1996 and 1997.) However, the combination is optional. For instance, it is perfectly possible to accept an inferentialist meta-semantics in combination with a truth-conditional semantics.

^{xii} It is important to note, however, that the connection with analyticity only holds if one accepts the assumption that there is a strict distinction between sentences that are meaning constitutive and those that are not, in accordance with (UAI) and (UAt). An alternative is to appeal to holistic understanding-assent links in one's semantics, along the lines of the cluster theory. Hence, it is possible to defend a version of the *semantic* UA-thesis (and not just the meta-semantic one) without committing to analyticity.

^{xiii} Williamson also tells the story of a second speaker, Stephen, who rejects 'Every vixen is a vixen', since he takes 'vixen' to be a vague term and believes that sentences containing vague terms have truth value gaps (2002:87).

^{xiv} In contradistinction to Eklund, for instance, who suggests that when two speakers dispute the status of a logical law they do both have the same competence dispositions. That is, the disagreement is explained as a result of one speaker having a blocked disposition (2002: 262).

^{xv} Although to what extent this is true depends on the case. It may well be that certain disagreements when it comes to the logical constants have more alarming consequences than Williamson recognizes. See Boghossian (forthcoming) for a discussion of one such case.

^{xvi} See Glüer 2003 for a discussion of the implications of Quine's belief holism for implicit definition accounts of meaning. As Glüer stresses, the challenge posed by belief holism is to come up with a principled distinction between meaning-constitutive and other sentences (2003: 52).

^{xvii} It is not plausible, for instance, to suggest that full understanding requires the ability to *explicate* one's concepts. This is a meta-level ability whereas the notion of understanding that we are after is a first-order ability (see XXX for a discussion).

^{xviii} See XXX.

^{xix} There is a reason for this, of course, since Williamson is primarily interested in rejecting epistemological analyticity which, again, depends on the truth of (UAI) and (UAt). However, since Williamson suggests that we can dispose of all forms of (UA)-theories, it is worth considering what the alternatives are.

^{xx} As Williamson would grant, since he explicitly denies that meaning is necessarily determined by the social practice (2007: 125).

^{xxi} See Boghossian (forthcoming) for a similar complaint.

^{xxii} Burge 1986 presents an argument to show that the experts themselves can be wrong. However, this argument is employed to show that meaning is *not* determined by the community practice. Hence, the version of externalism defended is no longer social externalism, as Burge explicitly recognizes.

^{xxiii} 1979: 90-91. Notice that the same is true in the case of Putnam. To have *acquired* a word, such as 'tiger', one must not only be situated in the right context, one's use must also 'pass muster' (1975: 248). For one's use to pass muster, Putnam suggests, one must have a set of 'stereotypical', well entrenched, beliefs commonly associated with the term (tigers are striped, carnivorous, etc). This, he argues, is required for mutual understanding and meaningful communication.

^{xxiv} This, of course, is Davidson's view. For instance, discussing the case of the speaker who utters 'There's a hippopotamus in my refrigerator', Davidson suggests that we may be right to interpret him as having said that there is a hippopotamus in the refrigerator. However, if *P* goes on saying that the hippopotamus is round with wrinkled skin and makes delicious juice, Davidson argues, 'we slip over the line where it is plausible or even possible to say correctly that he said that there was a hippopotamus in the refrigerator' (Davidson 1984).

^{xxv} Williamson's appeal to externalism seems, in part, motivated by this fear. It is an important constraint on a theory of linguistic meanings, he argues, that they can be shared across differences in belief and externalism allows for this (2007: 97).

^{xxvi} For a discussion of this see Pagin 1997. Of course, if the function is many-one, rather than an equivalence relation, a proper reduction of meaning facts to meaning determining facts will be ruled out. However, a supervenience relation suffices for the theory to be informative: it will tell us what meaning supervenes on, i.e. what the meaning determining facts are, and how the function from these facts to meanings (contents) is to be construed.