

Semantic Deference and Groundedness

ANTONIN THUNS

Centre de recherche en linguistique LaDisco, Université Libre de Bruxelles, Belgium

Semantic deference allows for the meaning of a word w a speaker uses to be determined by the way other speakers would understand or use w . That semantic deference has some role to play in semantic content attributions is intuitive enough. Nevertheless, the exact conditions under which semantic deference takes place are still open for discussion. A key issue that the article critically examines is Recanati's requirement that deferential uses be grounded, that is, that deferential uses be linked to non-deferential uses (Recanati 1997; 2000). After distinguishing between semantic and epistemic deference, I submit that the only way to maintain the Groundedness Thesis for truly semantic deference is to allow deference to idealized future linguistic collectives. I conclude that this is too high a price to pay for Groundedness and I suggest that it should be rejected as a semantic thesis.

Keywords: Semantic deference, word meaning, grounding.

1. *Semantic deference and the "Groundedness Thesis"*

Speakers defer semantically when the meaning of words and expressions they use is determined by other speakers' understanding of these words/expressions rather than by their own. Typically, when speakers do not master the semantic content associated with a word w they want to use, they defer to an external norm for the fixation of the meaning of w . The deferential mechanism is meant to account for the important fact that speakers can express meaningful thoughts involving lexical items they have an incomplete or imperfect understanding of. They do this by relying on the lexical competence and knowledge of other speakers.

Semantic deference seems to be supported by fairly robust intuitions about truth-conditions. It makes content ascriptions possible independently of speakers' lexical mastery and also represents the implicit commitments which go along with the use of a public language. It is indeed very natural to attribute its standard meaning to a word w

a speaker S utters, even if S's understanding of *w* is such that S would not be judged competent with *w* according to communal standards. If you overhear someone saying that 'The White House is threatened with subpoena', the most natural assumption is that the truth-conditions of this random speaker's assertion do not so much depend on her private understanding of the term 'subpoena' as on the way legal experts would understand it. By default, you assume that the speaker defers semantically to legal experts for the fixation of the meaning of 'subpoena'.¹ The fact that the speaker could misunderstand this technical expression does not make the truth-conditions of her utterance different from the truth-conditions the same sentence would have if it was uttered by you, or by a legal expert, in the same circumstances. Exactly like you or the legal expert, the casual speaker is entitled to *mean* subpoena by 'subpoena', in the double sense that she can not only *aim at* the relevant technical notion but effectively *summon* it in her discourse. This remarkable feat is supposedly achieved through semantic deference.

As Marconi observes, semantic deference is "real enough", and simply corresponds to the fact that "we do not regard ourselves as semantic dictators, like Lewis Carroll's Humpty Dumpty" (Marconi 1997: 90). As members of a linguistic community, speakers implicitly know that a lot more is going on about word meanings than what their limited perspective gives them access to. Language users are aware of the fact that meanings are not entirely up to them. Empirically, this is evidenced by speakers' willingness to have their language use corrected. It is also attested by a tendency to hold speakers accountable for the meaning of words they use, without regard for what they have in mind when using them. If I say there are 'courgettes' in the fridge, I make a false assertion if there happen to be only cucumbers there, even though it is precisely *these* cucumbers I had in mind (I bought them and put them there, I just tend to get confused about the labels). I could argue all I want, I asserted that there were courgettes in the fridge, not cucumbers.

So, however speculative the notion might appear at first sight, strong intuitions support the existence of a mechanism of semantic deference. These can be gathered both from the internal perspective of speakers and from the external perspective of hearers and other "content ascribers". Speakers are disposed to be corrected about word meanings and, even more interestingly, are disposed to recognize that, by the mere fact of using certain words, they express semantic contents they might not have internalized. Accordingly, hearers or other external observers standardly ascribe semantic content to words a speaker utters without regard for the possible idiosyncrasies in the speaker's lexical-semantic representations.

Clearly, the theory of semantic deference meshes with the "anti-individualist revolution" (Bochner 2014) of the 1970s, whose deep lesson

¹ From here on, the formula "S defers to *x* for *w*" will be used as a shorthand for "S defers to *x* for the fixation of the meaning of the word/expression *w*".

was that individual mental states do not always determine linguistic reference. Though Putnam does not use the term ‘deference’ in “The meaning of ‘meaning’” (1975), it is quite obvious that his famous point about the “division of linguistic labor” supports the idea of semantic deference. Like Putnam, you might be unable to tell the difference between elms and beeches, but the mere existence of the two labels ‘elm’ and ‘beech’ is evidence enough for you to suppose that two distinct tree concepts are involved, which more expert speakers master and are able to apply correctly. The same is true for ‘gold’, ‘molybdenum’ or even, as Putnam argued, for a perfectly common word like ‘water’, whose real extension-determining “meaning” H_2O is best defined by chemists and thus reaches beyond the superficial understanding that ordinary speakers might have of the word independently of scientific knowledge.² Burge (1979) made a similar theoretical point about social kinds such as ‘sofa’ and ‘contract’: these terms’ meanings depend on the socio-linguistic environment, not on the private conceptions of individual speakers. Burge’s “anti-individualism” also provides support for the idea of semantic deference. Generally, semantic deference fits in nicely within an externalist framework. Even better, it tends to strengthen such a framework by showing that ordinary speakers themselves are practical externalists. They seem indeed to go by the assumption that a great deal of semantic knowledge can be safely “outsourced”. In the division of linguistic labor we trust.

As appears from these considerations, semantic deference is tied analytically to the notions of knowledge and expertise. If speakers defer in the first place, it is because other speakers have a better grasp of the meaning of some terms and are experts at using them. Expert speakers provide standards of lexical competence. They are needed for the mechanism of semantic deference to work. This requirement is made explicit in Recanati’s formalization of the deferential mechanism. Recanati (1997) introduced a “deferential operator” which allows construing “deferential concepts”³ in a way analogous to indexicals. The deferential operator “ $R_x()$ ” is a function that applies to a “public symbol o ” (the word which requires deference to take place) whose content is that of the concept possessed by a “cognitive agent x ” (the expert) and yields a deferential concept “ $R_x(o)$ ” (the one available for the deferrer to use). In Kaplanian terms, the content of the deferential concept thus obtained is the same as the content of the non-deferential concept (the expert’s concept), but its character is different. The character of the deferential concept is metalinguistic. Informally, it is something like “whatever x means by o ”. So, when the random speaker uses the word

² This claim needs to be qualified. For Putnam, it is ultimately water itself, the actual stuff, and not chemists, that determines the meaning of ‘water’. This important point will be discussed below.

³ The use of the term ‘concept’ by Recanati might be somewhat puzzling. As I take it, ‘concept’ here refers to the general meaning of a term as it is understood by a particular speaker, that is, the *character* this term has for the speaker.

‘subpoena’, she implicitly points toward experts in the linguistic community and appeals to their semantic competence. For a deferential speaker, then, the character of the word ‘subpoena’ should be something like “whatever ‘subpoena’ means for the relevant expert”. But its content (i.e., its Kaplanian “intension”) is simply the content it has for the expert.

Experts are supposed to be part of the deferential speakers’ environment. As such, they are part of the linguistic context of deferential utterances. The x variable in the formula can be taken to be a contribution of the context to the utterance content – hence the analogy with indexicals. Deferential cases are cases of “social indexicality” (Recanati 2001). If there is nobody in the linguistic community to defer to, or if the “experts” themselves do not really understand what a term means, the deferential concept fails to acquire a determinate semantic content, very much like a demonstrative fails to refer in the absence of a demonstrated object. This means that a deferential concept can have a definite character but no definite content or intension, if it appears that in fact no one is in a position to use the term non-deferentially. Woodfield (2000: 442) provides the example of members of a religious sect who believe that “Jesus Christ and John Lennon are *alpha-enlightened* beings”, in a situation where no one (not even the guru) is able to really understand the phrase ‘alpha-enlightened’. In this case, the deferential operator applies to ‘alpha-enlightened’, but the deferential mechanism fails to yield any determinate (truth-evaluable) content, because no genuine expert is ultimately available to provide an interpretation for ‘alpha-enlightened’. This means that, in the actual world, there is no context that would endow the term with evaluable content. As a result, the assertion of the sect’s belief is semantically ill-formed and its truth-conditions cannot be calculated. By contrast, a student who acquires from her teacher a belief that “Cicero’s prose is full of ‘synecdoches” (Recanati 1997: 86) comes to entertain a functioning deferential concept of ‘synecdoche’, whose content can be traced back to a real expert.

Such a constraint on what counts as successful semantic deference constitutes what Recanati calls the “Groundedness Thesis”:

(Groundedness Thesis)

A deferential use is grounded only if someone at the other end of the deferential chain uses the expression in a non-deferential manner. (Recanati 2000: 452)

Deferential uses, Recanati claims, are parasitic on non-deferential uses. This means that deference must stop at some point. If everybody uses a word deferentially, then there is no real content associated with the word. It remains empty, meaningless. The Groundedness requisite can then be used to set apart functioning deferential uses from cases in which everybody defers to an authoritative interpretation which is in fact nowhere to be found. The link between semantic deference and se-

mantic expertise, together with the Groundedness Thesis (henceforth, GT) it leads to, constitutes the most natural account of the phenomenon under consideration, which I will refer to as the “standard account”. While I think the idea motivating the account is initially sound – because standards of lexical expertise do exist and can be identified – it raises a number of problems. These problems, I argue, call for severe revisions of the notion of Groundedness that end up making it unappealing.

2. *Epistemic deference, semantic deference and collective grounding*

The main problem of the standard account, Woodfield (2000) argues, is that it equates semantic deference with *meaning borrowing* from particular individuals. The application of the deferential operator to an expression σ seems indeed to allow the deferential speaker to “borrow” the meaning σ has for a particular expert (the one who happens to be related to the deferrer by a deferential chain). But what if the expert gets it wrong? And what if experts disagree among themselves? As Woodfield shows, this can lead to highly counter-intuitive misunderstandings between speakers of the same language, once several experts are involved in a deferential process. In the synecdoche example adapted from Recanati (1997), it is enough to imagine the following dialogue between Alf, the boy who picked up the word ‘synecdoche’ from his schoolteacher (the initial expert), and a linguist L (the new expert):

- (i) Alf says: ‘Cicero’s prose is full of synecdoches.’
- (ii) L replies: ‘No it is not. It’s true that his prose is full of figures of speech. But very few of them are synecdoches.’
- (iii) Alf replies: ‘I accept what you say. Cicero’s prose is not full of synecdoches.’

(Woodfield 2000: 448)

Let’s assume, with Woodfield, that the teacher to whom Alf defers himself has an imperfect understanding of the concept of synecdoche: he uses the term to refer to metonymies (of which synecdoches normally form a sub-kind). If we take the deferential operator to allow Alf to borrow the schoolteacher’s concept, then, when he says that ‘Cicero’s prose is full of synecdoches’, Alf is in fact asserting that Cicero’s prose is full of *metonymies*, since the operator forces us to treat the teacher’s understanding of the term as providing the semantic content of the deferential concept. But, in that case, the expert linguist L, who masters the meaning of ‘synecdoche’ and challenges the student’s assertion, would be in fact missing the point entirely. L is referring to synecdoches, whereas Alf, unbeknownst to himself, is referring to another notion that he borrowed from his teacher under the misleading label ‘synecdoche’. Since they are not co-referring, Alf and L are just talking at cross-purposes. Things become even more troublesome if we consider

the possibility that Alf's deferential pointer changed targets between (i) and (iii). If, after accepting L's judgement, Alf immediately stops deferring to the schoolteacher and starts deferring to L instead, then (i) and (iii) no longer stand in logical contradiction to one another, discursive appearances notwithstanding. At the level of propositional contents, the dialogue would indeed run as follows:

- (i) Cicero's prose is full of *metonymies*.
- (ii) No it is not. It's true that his prose is full of figures of speech. But very few of them are *synecdoches*.
- (iii) I accept what you say. Cicero's prose is not full of *synecdoches*.

Obviously, this misrepresents the communicative exchange, as it runs against our strong intuition that Alf and L do mean the same thing all along when they use the term 'synecdoche' (assuming they are engaged in a cooperative linguistic exchange), even if they do not understand it the same way and even if Alf's deferential attitude changes in the course of the exchange. Furthermore, it misrepresents the fact that Alf has apparently changed his mind about one and the same topic between (i) and (iii), namely, about synecdoches in Cicero's prose.

Fortunately, this problem can be overcome by drawing, with De Brabanter et al. (2005), a distinction between "semantic deference" and "epistemic deference".⁴ *Epistemic deference* is about directly exploiting others' cognitive resources. We defer epistemically to other people when we assume that they have a better understanding of the reasons and evidence for making certain claims we make. An agent A_1 who defers epistemically to another agent A_2 uses A_2 's words and behavior as reasons to make certain claims and draw certain conclusions, without necessarily understanding the criteria and reasonings motivating A_2 's claims and behavior. Epistemic deference is not a specifically linguistic phenomenon. As Rauti (2012: 326) points out, you may defer epistemically to a silent hunter in order to know which way to turn in a hunt for deer. Whether the hunter comments or not on the signs she spotted is irrelevant to your ability to defer epistemically to her. You do it as long as you make your own a claim or a belief based on someone else's knowledge or expertise, without thereby necessarily acquiring this agent's knowledge or expertise.

Semantic deference, by contrast, is a specifically linguistic phenomenon. It concerns the fixation of word meanings themselves, i.e., the fixation of their stable semantic contribution. A speaker S_1 may defer epistemically to another speaker S_2 for the reasons to use a word in a

⁴ In their taxonomic study, De Brabanter et al. (2005) clarify the distinction not only between semantic and epistemic deference, but also, within semantic deference, between default and deliberate deference. They also show that the problem of imperfect mastery must be considered separately, as it is not a necessary feature of deferential uses (one can deliberately defer to someone else's use/understanding of a word even in the absence of imperfect mastery, say, for playful or ironic purposes). I refer the reader to their work for further clarifications.

certain way. But it does not entail that S_1 defers *semantically* to S_2 for the very meaning of that word. As De Brabanter et al. (2005) show with their adaptation of Burge's arthritis example, a patient can defer epistemically to her doctor concerning the applicability of the medical term 'arthritis' to her condition. But she does not thereby defer semantically to this particular doctor for the fixation of the meaning of the word 'arthritis' itself. Even more, she cannot reasonably be expected to defer semantically to him.

Let's imagine that she is correctly diagnosed with arthritis by her physician, but that it appears later that he actually has a poor understanding of the ailment and, though able to diagnose it correctly (because he is good at identifying its symptoms), is not a reliable theoretical guide to arthritis (he holds many false beliefs about the ailment). Would the patient be willing to admit she was deferring all along and in every regard to her doctor for the *meaning* of 'arthritis', say, each time she was talking to people about her condition? We have reasons to doubt it. Even if her doctor was her only source of information about arthritis, it does not seem to imply she meant everything her doctor means by 'arthritis' each time she used the term. What she meant to refer to was the ailment that is supposed to be described by modern medicine under the label 'arthritis'. Consequently, it is more reasonable to assume that the patient, *by default*, defers semantically to the community of medical experts, and only defers epistemically to her doctor for practical purposes.

Coming back to the synecdoche example, it is easy to see that Alf and the expert linguist both defer semantically to the community of experts in rhetoric. They are definitely talking about the same thing: the semantic contribution of the word 'synecdoche' remains stable over the course of the dialogue. The difference between Alf's use and the linguist's use of 'synecdoche' is purely epistemic. Alf starts off deferring epistemically, first to his teacher, then to the linguist. He takes their claims as reasons to form particular beliefs involving the 'synecdoche' concept. The linguist, here, does not defer epistemically to anyone. But it does not mean that she would never revise her understanding of 'synecdoche', e.g., in the light of new developments in the field of rhetoric. She is still *disposed to defer*. So, if one accepts that default semantic deference consists precisely in this default disposition to defer, both Alf and the linguist defer semantically, by default, to the community of experts in rhetoric, which provides the ultimate norm regarding the correct use of the word 'synecdoche'.⁵

⁵ To be fair, this move is already implicit in Recanati (1997: 85), since he also writes that "there is a public interpretation of 'synecdoche', on which experts in rhetorics converge". What we have seen is that this conflicts with his other claim that "the symbol's content, in this context, is the content which the symbol 'synecdoche' has for the teacher" (92) because the teacher, as an isolated expert, could very well be mistaken about the norm.

Of course, epistemic deference may also have linguistic implications. If A_1 uses A_2 's *words* as a ground for epistemic deference, it may have an impact on the way A_1 conceives of the *meaning* of some of A_2 's words. If someone points out a tree to you and tells you 'That's an elm', you may or not decide to defer epistemically to that person. If you do, you make that person's claim your own and you can use it to tell someone else, pointing to the same tree, that 'That's an elm', or even to tell anyone else later that 'There was an elm there'. Obviously, this process might help you get a better grasp of the meaning of the word 'elm'. Likewise, Burge's Bert, the man who believes that arthritis can affect soft tissues as well as the joints (Bert believes that arthritis has lodged in his thigh), comes to a better understanding of the meaning of 'arthritis' once he is "apprised of the fact that arthritis is an inflammation of joints and cannot occur in the thigh" (Burge 1979: 198). Arguably, Bert may have come to this improved understanding of the meaning of 'arthritis' by deferring epistemically to other people's judgment involving the word 'arthritis', i.e., by making their claims his own without necessarily understanding the rationale behind them. So, in both cases ('elm' and 'arthritis'), it appears that a process of epistemic deference can have consequences on the deferential agent's lexical representations. But, as the discussion of the aforementioned examples is meant to show, it does not follow that epistemic deference has any role to play in the *fixation* of word meanings themselves. From the point of view of semantic content attributions based on deferential uses, local processes of epistemic deference are an idle wheel. Only *semantic-deferential* uses are the legitimate objects of GT.

As Woodfield notes, it is true that real-life processes of deference may stop at particular experts, but it is not true that they have to. An expert A at the end of a deferential chain could always herself defer for some reason to other experts B and C, while B and C defer to A for different reasons (Woodfield takes the example of scientists working together on a disease). Even if A stops deferring to B and C, another expert D can always show up with new evidence and reconfigure the deferential relations. It is then possible to imagine never-ending "*circles* of mutually deferring agents" (Woodfield 2000: 435). The result is that no one is in a position to use a given technical term in a completely non-deferential way, which contradicts GT.

This should not be a problem. Once semantic deference to the "first expert at hand" (De Brabanter et al. 2005: 14) is abandoned and the distinction between epistemic and semantic deference is introduced, the idea of mutually deferential collectives starts to make a lot of sense. Real-life experts only provide partial and fallible ways to approach the meaning of a word. Laypeople defer to them epistemically, not semantically. Likewise, experts may defer epistemically to one another in various ways, but the existence of ideal *individual* experts, worthy of full semantic deference, is highly implausible and maybe simply not required for *semantic* purposes.

The conclusion of this section is that, if semantic-deferential uses have to be grounded, the only plausible “cognitive agents” for the deferential operator to index are *linguistic collectives* made up of variously competent and mutually deferential speakers. This, it seems, follows naturally from the distinction between epistemic and semantic deference. Now, this requires making the notion of a linguistic collective precise and fit to play the role of the “cognitive agent” demanded by the deferential operator. What, exactly, is the “collective” target of the deferential operator? Two options are open. The operator could be targeting only the best experts of the field. Even if none of them individually possesses all the available knowledge, their combined knowledge might deliver the best possible definition of a term. Alternatively, the operator could be targeting something more abstract: a common body of knowledge, a set of norms or even a public language taken as a whole.⁶ Perhaps the second option fits better with our intuition that common knowledge is more than the mere sum of individuals’ knowledge and is exempt from idiosyncrasies. Moreover, parts of it could even not be known by anybody but nevertheless available, stored in books or, as is increasingly the case, in computers.

To be sure, the notion of an abstract linguistic standard is hard to pin down and might raise ontological questions.⁷ Committing to such a notion, however, seems to be a first price to pay if one wants to maintain GT for truly *semantic* deference (as opposed to epistemic deference). Yet, as I am about to argue, even clarifying this notion would not get proponents of GT out of trouble. Indeed, in order to maintain GT, one must make sure that collectives *in fact* play the role of higher-order experts stabilizing meanings. So far, the discussion of ‘arthritis’ and ‘synecdoche’ only shows that we are inclined to think that they *should* do so. Yet it does not mean that they do. As the next section purports to show, the strong ties between semantic deference and externalism indicate that they do not.

⁶ All these suggestions are found in De Brabanter et al. (2005): “[The] truth value [of an utterance of ‘I have arthritis’] is determined by appealing to the experts, and to the linguistic community more generally, regarding the question of what counts as ‘arthritis’” (4); “In English, the meaning of ‘arthritis’ is established in connection with the common body of medical knowledge” (5); “[The arthritis patient] is deferring by default to the norms of the linguistic community” (9); “we will be using the variable *x* either for users of a language or for the language itself” (8).

⁷ As O’Madagain (2014) points out, in order for the conventional semantic value of a term to be determined by a group, groups must be susceptible of being attributed intentional states and concepts. However, the model of group semantics he offers, inspired by that of Lewis (1969) for conventions, only works for attributions of intentional states to groups that are the result of an explicit procedure (as is the definition of the term ‘meter’ by the International Bureau of Weights and Measurements). This account thus falls short of explaining how distributed knowledge could ground the semantics of terms whose meaning is not fully determined by explicit agreement.

3. *Deferring beyond current understanding: virtual grounding*

The deep problem with the “standard account” of semantic deference, even in the amended version where abstract collectives replace individuals, is that it rests on descriptivist assumptions. Saying that speakers defer to experts’ *understanding* of a term is tantamount to saying that speakers defer to the experts’ *definition* of that term. Once we accomplish the move to abstract linguistic standards, the underlying assumption is that the community provides *the most complete definition*, while no individual expert speaker possesses all the relevant information. Yet, at least in the case of natural kind terms, descriptivism is usually held to be untenable, because questions about a term’s meaning can hardly be separated from questions about our epistemic access to the term’s denotata. Given our epistemic fallibility, the semantic import of natural-kind term is taken to reach beyond available descriptions of the denotata and, even, to be independent of them. What, indeed, if the experts collectively got some facts wrong concerning the referents of natural kind terms, so that no appeal to collective wisdom can compensate for individuals’ shortcomings? Can collective misconceptions ground meanings in the desired sense? Even if collective conceptions are not entirely mistaken, they might simply fall short of providing the determinacy of meaning that we take natural kind terms to have. And how are we to deal with radical changes of conception that occur in the course of a term’s history (for example, the fact that fire is now thought of as the effect of a chemical reaction and no longer as a fundamental element)? These sorts of questions lead, from the rather uncontroversial observations that speakers defer and that there is division of linguistic labor, to more radical claims about meaning externalism.

As Putnam famously argued, the fact that our conceptions of natural kinds change does not entail that the *meaning* of natural kind terms changes accordingly. His point was precisely that the meaning of natural kind terms depends on the external world itself, independently of anyone’s grasp of the corresponding referents, hence independently of available *descriptions* of those referents. As Liu (2002) highlights, Putnam’s thought experiments about natural kind terms and Twin-Earth aim above all to establish the doctrine of *physical externalism*, according to which the meanings of natural kind terms are determined by objective natural boundaries, independently of the state of our collective knowledge. “[M]eanings’ just ain’t in the head!” (Putnam 1975: 227) implies that meanings are external to any head, society’s heads included. Even so-called experts do not determine meanings. At most, they only provide the best approximations of how meanings should be construed. Therefore, experts, even collectively, do not ultimately ground our use of natural kind terms like ‘elm’, ‘beech’, ‘water’ or ‘fire’. Only nature does. In the face of this, it is fallacious to jump from Putnam’s “division of linguistic labor” to the standard account of semantic

deference in which experts ground meanings. It is, for example, fallacious to attribute to Putnam, as I did above, the view that chemists determine the meaning of 'water'. Chemistry is at best the most refined method to approach the nature of that which we call 'water', but the meaning of the word, according to Putnam, depends on the kind of substance water actually happens to be, without regard for our conceptions. Twin-Earth-style thought experiments thus imply a rejection of meaning descriptivism for natural-kind terms.

If these externalistic intuitions about meaning are correct, then the conditions under which GT can be maintained must be amended further. Speakers must be taken to defer semantically *beyond current standards of expertise and understanding*. What, in this case, are they deferring to? The intuition behind physical externalism seems to be that speakers must be taken to defer to "the world", to nature itself and the substances it contains. Though it accounts for truth-conditional intuitions, this construal of semantic deference is a considerable departure from the original notion of deference and one might object that it does not deserve the label "deference" at all (De Brabanter and Leclercq, ms). However, I think that it is possible to help ourselves to Peirce's conception of knowledge (Peirce 1878) in order to recast this "semantic deference to the world" in a way that preserves the sociolinguistic nature of deference (speakers defer to other speakers) and the descriptivism underlying the reconstruction of the phenomenon. This can be done by allowing speakers to defer to *idealized future collectives of experts*. According to this interpretation, the meaning of the words 'fire' or 'water' is determined not by what current collectives of experts happen to agree on, but by what future collectives will discover about the "substances" in question.⁸ So, instead of a rather metaphorical deference to the world, one can maintain the basic notion of deference to other speakers, except that these now are hypothetical future speakers that we take, by stipulation, to be "ideal" experts, in possession of an understanding/definition so perfect that it is as good as the world itself to determine the extension-determining meaning required by our truth-conditional intuitions (those motivating semantic deference in the first place). This kind of *counterfactual* or *virtual deference* (deference to virtual ideal speakers) is a way of neutralizing the undesirable consequences of the descriptivism underlying talk of semantics deference whilst maintaining the structure of the standard account as well as GT. Obviously, since "cognitive agent *x*" is now located in the future, the amended account leads to the introduction of a notable *teleological* dimension in Groundedness. Paraphrasing Peirce's formula, one could summarize the account by claiming that the meaning of that which a word encodes does depend on the real fact that investigation

⁸ The proposal is a form of "Temporal Externalism" in the sense of Jackman (2005), except that it is augmented by the necessary idealization of future collectives (see below).

is destined to lead, at last, if continued long enough, to stabilizing that meaning.⁹

However, the Peircean solution only seems to apply to alleged natural-kind terms: a notion that is notoriously difficult to clarify. One might indeed object that words like ‘water’, ‘elm’, ‘fire’ and ‘arthritis’ all designate “natural kinds” only in a very broad and unenlightening sense of the term. As to physical externalism, it faces the “*Qua*-problem” (Devitt and Sterelny 1999: 79): How can we know that speakers, upon naming samples of a given entity that we *retrospectively* interpret with the categories of contemporary science (‘Water has a chemical formula’; ‘Elms form a genus’; ‘Fire is a chemical reaction’; etc.), actually converge on the implicit norm of the hidden nature affording the natural-kind interpretation? Whilst acknowledging the difficulty, I will not engage in any epistemological discussion about the grounding of natural kind terms here. As I now will argue, the point on “deference beyond current standards” is much more general and does not crucially depend on the outcome of such a discussion. It can even be made by considering terms with a far more dubious epistemological status than alleged natural kinds.

Take *polemical terms* such as ‘justice’, ‘science’, ‘philosophy’, ‘democracy’. They pose a specific problem for the standard deferential account, because their meaning is still so heavily debated in the community of experts, without any consensus emerging, that they suffer, as they stand, from a lack of determinacy which threatens the propositionality of the sentences that contain them, hence the possibility for semantic deference to be grounded. Most philosophical and moral terms are essentially polemical. Concepts debated within the social sciences and linguistics are also often polemical. Just consider, to take but one example, the issue of the distinction between what counts as ‘semantic’ and what counts as ‘pragmatic’. The range of conflicting views available is such that it could be argued, by a purely external observer endorsing the Groundedness principle, that these alleged terms of art are in fact deprived of meaning. I take it that a similar point could be made about the need to distinguish between ‘synecdoche’ and ‘metonymy’.

Yet, experts disagree precisely about what they take to be one and the same conceptual problem, which they represent by using the same expression. They consider that there should be only one version of the ‘semantic/pragmatic’ distinction,¹⁰ or one clear-cut definition of ‘syn-

⁹ “‘Truth crushed to earth shall rise again’, and the opinion which would finally result from investigation does not depend on how anybody may actually think. But the reality of that which is real does depend on the real fact that investigation is destined to lead, at last, if continued long enough, to a belief in it.” (Peirce 1878: 274)

¹⁰ This claim is compatible with the facts that experts currently recognize that there are many theoretically sound ways of making the distinction. I think that this is due to our epistemic limitations. At the end of the day, the presumption behind the use of dichotomic scientific labels is that there should be a clean way to “draw the line”. If there isn’t, then the dichotomy should be abandoned or reformulated.

ecdoche', because these are supposed to capture something about the investigated objects themselves. In the same vein, many philosophers assume that there is something substantial to our intuitions about 'morality' and 'justice', or about the autonomy of the 'political', and engage in passionate discussions about these topics, whereas the very existence of valid objects of inquiry has not been settled yet (and, in these cases at least, cannot be settled simply by pointing at natural samples of the referents, as physical externalism would have it). We can have the intuition of the unicity of an object/topic and coin a term for it long before we know whether the intuition is justified in the first place. In such cases, a deferential attitude appears to be almost constitutive of, and built into, the use of the term.¹¹ This means that there should be some kind of external linguistic standard to which speakers are taken to defer by default. Given the currently polemical nature of these terms, the linguistic standard that speakers are gesturing toward only comes as a promise, as the possibility of a future agreement.

Let's recapitulate. Whether we consider natural kind terms or polemical terms, the point is basically the same: their semantic content is not plausibly stabilized – hence not grounded – by current collectives of experts. Yet we have the intuition they should be grounded for the sake of semantic content attributions. Therefore, GT has to be amended in the proposed way: by allowing default semantic deference to idealized future collective states. Deference to idealized future states of the community reflects the situation in which, on the one hand, there is some implicit agreement that there should be a stable meaning behind a given term and, on the other hand, the community of experts is unable to ground it, either because the extension-determining meaning of the term depends on the hidden nature of its referent, which only a hypothetical perfect state of knowledge could fully reveal (natural-kind terms), or because its meaning depends on a hypothetical perfect coordination between future experts (polemical terms).

This argument, I believe, works for all terms whose meanings have a teleological component. In our proposed Peircean construal, this concerns natural-kind terms as well as polemical terms. *Teleological terms* refer to entities our understanding of which is always perfectible and which are used in a process of ever-growing knowledge (natural-kind terms) or ever-finer rational argumentation (polemical terms). Follow-

¹¹ Kaufmann (2006) explores the idea that "socio-political concepts" (expressed by terms like 'nation', 'God', 'public opinion') are constitutively deferential because they do not exist independently of the communicative chains through which users get acquainted with them. The way they are given is purely communicational. Ordinary agents gain access to them only by deferring to higher authorities (Church, political elites, social scientists) who also happen to use the terms deferentially. Experts defer to higher authorities or principles, which are only evidenced through language (writings, oral traditions, testimonies), and, crucially, never through the provision of an extra-linguistic referent. The deferential chain never stops: socio-political concepts only have a reference "on credit". As such, they blatantly violate the Groundedness principle in its original formulation.

ing Jackman (2005), one may *characterize* (in the technical sense of *character*) the meaning of teleological terms as being sensitive to “epistemic” considerations. This means that changes in our conceptions of the meaning of these terms are brought about by epistemic factors: new discoveries or a better organization of our knowledge. Jackman contrasts “epistemic” with “pragmatic” explanations of linguistic change:

“Pragmatic” theories explain change in terms of nonrational factors such as taboo, metaphor, semantic drift and the like. For instance, ‘zipper’ changes from being a brand name to a generic name for such mechanical fasteners because the brand is so successful that users of the language gradually forget that the items of that kind are ever called anything else. We have a sociological/psychological explanation of the change, but no justification in terms of the truth of the beliefs involved. By contrast, “epistemic” theories explain changes in usage in terms of factors such as the need to keep our beliefs consistent both with new experience and with each other. We stopped, say, applying the term ‘fish’ to whales because we discovered that whales were in many important respects closer to those creatures we called ‘mammals’ than to other creatures we called ‘fish’. (Jackman 2005: 370)

In Jackman’s view, we take epistemically-driven changes to already affect what we mean with teleological terms, whether we can foresee these changes or not (and whether speakers explicitly acknowledge being “temporal externalists” or not). In the Peircean version of temporal externalism I claim proponents of GT are committed to, laypeople, individual experts and currently existing linguistic collectives use teleological terms deferentially. They all defer beyond the limited horizon of the community’s current epistemic state relative to the concepts at stake. They defer to a more advanced state of the debate, to an ideal state of the epistemic community in which a critical mass of interpretations would converge. What matters for semantic content ascriptions is not the actual criteria speakers use, but the fact that they are collectively engaged in using a term with the presumption that it corresponds to a valid concept. As a result, the ultimate semantic import of the term is more virtual than actual and calls for a notion of *virtual semantic deference*, i.e., deference to a potential or future normative agreement on the meaning of the term. Committing to such a notion, I suggest, is the price to pay if GT is to be maintained regarding deferential uses of teleological terms.¹²

¹² Schroeter and Schroeter (2014) propose a “connectedness model” purporting to solve the question of the semantic content of “normative concepts” along very similar lines. Their model abandons the requirement of sameness of criteria and replaces it with a “tradition-based determination theory” of semantic values: the “entire representational tradition (*i.e.*, the entire set of token thought elements bound together by relations of apparent *de jure* sameness)” (2014: 12) is taken as the default semantic value for the word/concept. Historically extended representational traditions encompass conflicting views and could therefore be taken as the default value of what I called “polemical concepts”. If representational traditions are extended in the past, it seems that they also point toward the future: future epistemic states of the community and future agreements between experts might solve the current lack of referential anchoring of some traditions.

4. *Too tough a bullet to bite?*

I have argued that GT commits the theorist to far more than the existence of expert speakers, as long as the thesis is to apply to semantic deference and not merely to epistemic deference. Semantic deference concerns the ascription of truth-conditions to utterances. Given the implicit externalism of our linguistic practice concerning certain terms, truth-conditions are world-involving and cannot be fully determined by mental states alone, even by the mental states of the most knowledgeable speakers. Abstracting away from individual speakers towards the communal level fits better with what speakers take themselves to be doing, as the ‘arthritis’ example illustrates. Yet, upon closer scrutiny, deference to more abstract linguistic standards is not enough to vindicate our truth-conditional intuitions. As Jackman (2005) points out, changes in usage guided by epistemic factors (for example, the discovery that whales are not fishes) are normally not seen as changes of meaning, so that we also need to treat future “conceptual developments” as already affecting the meaning of terms we use. The fact that whales have been reclassified as mammals at some point does not seem to entail that the meaning of ‘whale’ has changed. We have simply acquired a better grasp of the properties that delineate the extension of the term, that is, a better grasp of the extension-determining meaning of the term.¹³ So, the intuition is that extension-determining meanings are already there, fully determined, despite the fact that we only very imperfectly apprehend them. Supposing the phenomenon of semantic deference is to retain its socio-linguistic dimension and be grounded at the same time, I have submitted that there is no other option for the deferentialist than to accept that we defer, implicitly and by default, to idealized future states of collective understanding.

Now, my final claim is that the amended account of Groundedness I have just sketched is probably too committal. For one thing, it obviously rests on a very severe idealization. But, much more problematically, *the account abusively attributes a foundational role to intuitions*. This last problem, I think, is serious enough that it is preferable to reject GT altogether and reframe the issue in less demanding terms.

It is one thing to recognize, as Jackman does (2005: 370), that, in the counterfactual situation where we could see into the future and discover an epistemically superior usage of one of our terms, we would be willing to change our own usage accordingly (thereby showing that we take ourselves to be same-meaning with future speakers). It is quite another to assume, as I have done for the sake of dramatization, that some metaphysical connection between our usage and a hypothetical future perfect understanding *presently* grounds the semantic import of teleological terms. Viewed in this light, virtual grounding is nothing

¹³ Of course, other, non-denotational aspects of the “meaning” of the term, such as prototypical representations and stereotypes, may have changed as a consequence of the reclassification. But these are not supposed to affect literal truth-conditions.

more than “temporally loaded”¹⁴ Platonism. Yet I claim that this is precisely what GT inevitably leads to, if *semantic* deference is supposed to lead to non-deferential uses.

Perhaps one way to vindicate the intuitions motivating talk of semantic deference whilst avoiding outlandish metaphysical claims is to reject Groundedness as a metaphysical thesis and recast it as a psycholinguistic hypothesis. Put very briefly, the rationale behind this move would be the following. The external world is essentially inscrutable. Likewise, our relation to the external world is inscrutable. As a consequence, the extent to which the use of our terms follows precise epistemic norms (e.g., the norm of the “hidden nature”) is not known a priori, essentialist intuitions notwithstanding. Admitting that, it is not inconceivable that a great deal of our terms whose meaning is sensitive to epistemic considerations (“teleological terms”) are in fact semantically indeterminate: their boundaries evolve with our dealings with the world without being firmly settled. There is no reason to assume, just because we have realistic semantic intuitions, that nature itself or future speakers are systematically cooperating to fully determine the semantic import of the teleological terms we use. However, this does not prevent us from *taking* these terms to have determinate contents, exactly in the same way that we take ourselves to speak English and not just similar idiolects. It is a psycholinguistic fact that teleological terms appear to us to be endowed with substantive content and that, most of the time, a change of our understanding does not appear to us as a change of meaning. In fact, it is possible that we do not know whether the extension-determining meaning of a given term fully obeys “epistemic” rather than “pragmatic” considerations. Consequently, our commitment to the stable semantic import of teleological terms could be, at least in a significant number of cases, the result of an illusion of the semantic kind.

Take the case of the word ‘whale’. Most likely, its extension-determining meaning has not changed over the last few centuries, because the term plausibly had the same extension in 18th century fishers’ mouths as it has in ours. But let’s imagine a counterfactual situation in which ancient fishers had come across some species of very large fish that looked very much like whales. I do not see any reason not to include this species in the extension of ‘whale’ in their English, if ancient fishers would so call those large fish. This would not be incompatible with the fact that subsequent speakers, once they have discovered better ways of classifying species, would perhaps say that ‘in fact those fishes weren’t whales’ – precisely the kind of argument the temporal externalist appeals to. To be sure, speakers could opt for the exclusion of large fish from the extension of ‘whale’, thereby abiding by the scientific use. But it does not follow from this that whale-like fish were not in the extension of ‘whale’ before the discovery of the fact

¹⁴ Jackman’s phrase.

that *most* of the animals called ‘whales’ are mammals. The extension of the term has simply been changed. Therefore, the informal claim that whale-like fish were incorrectly called ‘whales’ is a *retrospective semantic illusion*. In this scenario, whale-like fish just *were* in the extension of ‘whale’. They simply no longer are. Alternatively, in the same counterfactual situation, speakers might just refuse to follow epistemic considerations and keep using the word ‘whale’ as they used to (i.e., as a term referring to both whales and whale-like fish), leaving to biologists the coinage of other terms for the purpose of precise classification. In this second scenario, whale-like fish unproblematically were and remained in the extension of ‘whale’. In fact, something very similar happened in the actual world for the term ‘reptile’. Scientifically, this term is no longer deemed to coincide with a consistent natural kind. In ordinary language, birds do not count as ‘reptiles’, whereas crocodiles, turtles, snakes and lizards do. In cladistic terms, this does not make much sense, for crocodiles are phylogenetically closer to birds than to the other ‘reptiles’. In spite of this discovery, ‘reptile’ has remained a perfectly serviceable English word and its extension has remained the prescientific one. Theoretically, its use could have been extended to include birds, and one could have said that ‘in fact, birds are reptiles too’, but that is not how usage has evolved. Instead, there is still an ordinary use for ‘reptile’ (probably motivated by the superficial properties so-called reptiles share, which provide useful contrast with birds, mammals and amphibians, that is, motivated by “pragmatic” factors) and no well-founded scientific use.

I take it that such examples are easy to multiply, showing that externalist intuitions about the “true scientific extension” of a natural-kind term suffer from what one might call a “scientific-essentialist” bias. My point is not that usage is never sensitive to essentialist considerations. It is rather that, even if ordinary usage does end up following scientific norms, it is only a retrospective illusion that the term “always had” the extension-determining meaning and the extension that we now take it to have. A less committal thesis would be that such illusions and other objectivistic intuitions about word meanings are compatible with the *de facto* semantic indeterminacy of a number of teleological terms and the open-ended nature of the semantic deference attached to them. The intuitions on which talk of deference and groundedness is based may have an important psychological role to play and are probably indispensable to the project of knowing the world and communicating about this knowledge. However, by themselves, these intuitions are not enough to *ground* semantic deference, except if one is disposed to embrace a form of teleological Platonism. Unless one is willing to bite that bullet, I take it that GT should be rejected.

5. Conclusion

Once linguistic deference is taken to determine *semantic* content proper, it forces us to go beyond the level of concrete interactions, or even chains of interactions, between speakers and to consider more abstract linguistic standards that serve as guiding principles for our semantic content ascriptions. More precisely, only virtual semantic deference to future idealized linguistic collectives seems to be able to secure the intended semantic import for our teleological terms. Barring the explicit endorsement of a form of Platonism, such an account is too committal and is better recast as a psychological hypothesis.

My focus has been mainly on prototypical examples of semantic deference: deference for terms like ‘elm’, ‘water’, ‘arthritis’ and ‘synecdoche’, all amenable to scientific enquiry. I also addressed the category of polemical terms, which motivates the introduction of the more inclusive category of teleological terms: terms pointing at some “realities” (not necessarily physical) which might never be fully known but to which speakers nonetheless take themselves to be referring through words. I take the category of teleological terms to encompass all the terms whose meaning is sensitive to epistemic considerations: natural kind terms, scientific terms, polemical terms.

Teleological terms provide prime illustrations of semantic deference, because examples involving our epistemic limitations are the ones that come to mind most easily. A question worth exploring, though, is whether speakers can defer semantically for words whose meaning does not fit nicely in the expertise-based model that underlies the discussion here. Besides teleological words, there is a vast grey area of relatively imprecise words, which are nevertheless easily grasped by speakers because of their familiarity. For most of these perfectly ordinary items, it is not clear that it is possible to identify a community of experts, nor a teleological dimension attached to them. Are there specialized bodies within the English-speaking community who stabilize the lexical meaning of the words ‘meeting’, ‘sadness’, ‘soft’ or ‘friend’? Is the linguistic community engaged in the pursuit of a better understanding of the semantic import of these terms? Most likely not. Moreover, it is arguably part of the very functions of ordinary non-technical words to remain semantically underdetermined, in order to provide very abstract and flexible frames for the pragmatics of everyday communication to fill in.

Yet these words intuitively express distinct and graspable notions. Even though they are non-technical, vague and extremely context-sensitive, they are not obscurely indeterminate in the sense esoteric concepts are. So, it seems that someone who does not understand an ordinary word like ‘friend’ (a small child, typically, or a foreigner), could defer semantically to other speakers for such a term. Would this kind of deference be constrained by GT? What would it mean to say that another speaker uses the word ‘friend’ non-deferentially? Since

the general concept of ‘friend’ is semantically underdetermined, it can be used with divergent semantic imports by equally competent speakers. Does it mean that each non-deferential speaker grounds ‘friend’ in their own fashion, so that the concept is “multiply grounded”? This is an interesting possibility, but it runs counter to our understanding of the implicit motivations of the deferential speaker. Someone who defers semantically for the meaning of ‘friend’ does not want to wind up expressing any old speaker’s idiosyncratic take on the word. The deferential speaker, by default, aims at the standard, collective concept. But what is the collective concept of ‘friend’ and how is it grounded? It does not seem to be a teleological term, because its use is more sensitive to pragmatic rather than epistemic considerations. Knowing that future speakers will use the term differently would not be a reason for us to consider that our current use of ‘friend’ is mistaken in any way. And yet, what if future psychology converted words like ‘friend’, ‘love’, ‘trust’ or ‘doubt’ into full-blown natural-kind terms, so that semantic appropriateness of past uses could be reevaluated with respect to new scientific findings?

These reflections are, of course, purely speculative, but the fact that they make sense hints at the positive side of Section 4’s conclusion. If, on the one hand, it is possible that many of our teleological terms are semantically indeterminate, it is also possible, on the other hand, that some apparently non-teleological terms are more determinate than we anticipate, because they could also end up falling under epistemic considerations. In a sense, the whole project of the “special sciences” (sociology, psychology, anthropology...) is to unveil the “hidden” social and psychological functions of a host of ordinary terms, i.e. to uncover their externalistic semantics. The fact that the terminology of the special sciences is “contaminated” by ordinary language and always runs the risk of being semantically indeterminate is just a particularly dramatic case of the general risk of semantic indeterminacy affecting all our terms, natural-kind terms included. Once GT is rejected as a semantic thesis, strict boundaries between kinds of terms are no longer needed. The pervasiveness of semantic deference and its psychological importance can be fully acknowledged without us committing to a problematic metaphysics of meaning.*

References

- Bochner, G. 2014. “The anti-individualist revolution in the philosophy of language.” *Linguistics and Philosophy* 37 (2): 91–120.
- Brown, M. 2010. *How I Killed Pluto and Why It Had It Coming*. Spiegel & Grau.
- Burge, T. 1979. “Individualism and the mental.” In his *Foundations of Mind, Philosophical Essays, Volume 2* (2007). Oxford: Clarendon Press: 100–150.

- De Brabanter, P., Stojanovic, I., Villanueva Fernandez, N. and Nicolas, D. 2005. "Deferential Utterances." Retrieved from https://hal.archives-ouvertes.fr/ijn_00000575v1.
- De Brabanter, P. and Leclercq, B. Manuscript. "Semantic externalism(s) and semantic deference."
- Devitt, M., and Sterelny, K. 1999. *Language and Reality. An Introduction to the Philosophy of Language*. Cambridge: The MIT Press.
- Jackman, H. 2005. "Temporal externalism, deference, and our ordinary linguistic practice." *Pacific Philosophical Quarterly* 86: 365–380.
- Kaufmann, L. 2006. "Les voies de la déférence: Sur la nature des concepts sociopolitiques." *Langage et société* 117: 89–116.
- Lewis, D. 1969. *Convention*. Cambridge: Harvard University Press.
- Liu, J. 2002. "Physical externalism and social externalism: Are they really compatible?" *Journal of Philosophical Research* 27: 381–404.
- Marconi, D. 1997. *Lexical Competence*. Cambridge: The MIT Press.
- O'Madagain, C. 2014. "Can groups have concepts? Semantics for collective intentions." *Philosophical Issues* 24 (1): 347–363.
- Peirce, C. S. 1878. "How to make our ideas clear." In C. J. W. Kloesel (ed.). *Writings of Charles S. Peirce: A Chronological Edition. Volume 3 (1872-1878)* (1986). Indiana University Press: 257–276.
- Putnam, H. 1975. "The meaning of 'meaning.'" In his *Mind, Language and Reality. Philosophical Papers, Volume 2* (1975). Cambridge: Cambridge University Press: 215–271.
- Rauti, A. 2012. "Multiple groundings and deference." *The Philosophical Quarterly* 62 (247): 317–336.
- Recanati, F. 1997. "Can we believe what we do not understand?" *Mind and Language*, 12 (1): 84–100.
- Recanati, F. 2000. "Deferential concepts: a response to Woodfield." *Mind and Language* 15 (4): 452–464.
- Recanati, F. 2001. "Modes of presentation: perceptual v. deferential." In A. Newen, U. Nortmann and R. Stuhlmann-Laeisz (ed.). *Building on Frege: New Essays on Sense, Content, and Concept*. Stanford: CSLI Publications, Stanford University, 10: 197–208.
- Schroeter, L., and Schroeter, F. 2014. "Normative concepts: a connectedness model." *Philosophers' Imprint* 14 (25): 1–26.
- Woodfield, A. 2000. "Reference and deference." *Mind and Language* 15 (4): 433–451.

* My research is funded by the Fonds de la Recherche Scientifique (F.R.S.-FNRS) and is part of the research project *Cognitive transparency of semantic contents, pragmatic determination of reference* (PDR T.0184.16, 2016-2020). I would like to thank all members of the project for their very helpful comments on earlier versions of the article.