

Deferential Utterances

Isidora Stojanovic, Philippe De Brabanter, Neftali Villanueva Fernandez,
David Nicolas

► **To cite this version:**

| Isidora Stojanovic, Philippe De Brabanter, Neftali Villanueva Fernandez, David Nicolas. Deferential
| Utterances. 2005. <ijn_00000575>

HAL Id: ijn_00000575

https://jeannicod.ccsd.cnrs.fr/ijn_00000575

Submitted on 3 Feb 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deferential Utterances

PHILIPPE DE BRABANTER, NEFTALÍ VILLANUEVA FERNÁNDEZ, DAVID NICOLAS, AND ISIDORA STOJANOVIC¹

1 Introduction

Our aim in this paper is to clarify the distinctions and the relationships among several phenomena, each of which has certain characteristics of what is generally called “deference”. We distinguish linguistic deference, which concerns the use of language and the meaning of the words we use, from epistemic deference, which concerns our reasons and evidence for making the claims we make. In our in-depth study of linguistic deference, we distinguish two subcategories: default deference (roughly, the ubiquitous fact, noted by externalists like Burge or Putnam, that the truth conditions of our utterances are determined with respect to the language parameter supplied by the context), and deliberate deference (roughly, the intentional, communicative act of using a given expression the way it is used in some contextually specified idiolect or dialect). We also discuss the phenomenon of imperfect mastery, often as-

¹ The order of the names is not intended to reflect differences in the work done by each author.

sociated with deference, and which we show to be independent both of linguistic deference and of epistemic deference. If our analysis is correct, then some recent debates on deference (e.g. between Recanati and Woodfield) can be shown to result from a failure to appreciate all the distinctions that we draw here.

1.1 A quick look at the literature

The enthusiastic reception of Hilary Putnam's dictum –“meanings are not in the head”– triggered a wave of interest in certain kinds of language use that intuitively support Putnam's thesis. Externalists argued that the contents of our thoughts and utterances are not individuated by processes internal to the agent, but by the way the world happens to be. Tyler Burge (1979) offered examples to show that two distinct individuals could be, qualitatively, in the same internal state, and still have different beliefs. Particularly enlightening are examples involving imperfect mastery, that is, meaningful uses of concepts by speakers or thinkers who do not really master them. Such partial understanding, Burge argued, is omnipresent in the use of language. Brian Loar took the lessons from those cases of “falsity–due–to–misunderstanding” one step further, arguing that the satisfaction conditions of many concepts depend not on our personal epistemic abilities, but on language, “that social fact to which we defer” (1990, 118).

Two intuitions are at conflict here. On the one hand, we do not master all the concepts that we use. Still, we would not want to say that it is impossible for us to have or express thoughts about, say, rockets, aluminum or contracts just because our understanding of those concepts is incomplete. On the other hand, it is not clear how we manage to say anything true or false when we don't master the concepts associated with the words we are using. Talking nonsense is not the same as saying something false.

The idea of semantic deference was introduced precisely in an attempt to reconcile these conflicting intuitions. Its rationale is that we can meaningfully use concepts that we do not master because we defer to the public language, whose rules are fixed by our linguistic community.

François Recanati has postulated the presence of a “deferential operator” in all cases of deference. The deferential oper-

ator is an unarticulated constituent that affects the truth conditions of an utterance. It applies to any symbol (whether or not the symbol belongs to the language of the utterance) and delivers a well-formed expression, provided that there is tacit or explicit reference to some user of this symbol, the deferee (cf. Recanati 2000, 282, 272; 2002).

However, given that we do not master most of the concepts we use, does that mean that a deferential operator affects the semantics of most terms occurring in our utterances? Recanati thinks so: deference is present in every case of imperfect mastery, from the layman who uses technical terms in court to children learning their mother tongue.

This stance has raised criticism. Andrew Woodfield has argued that deference is not just something that happens to any language user. Deferring, he says, is an “intentional act, done by a person for a reason” (Woodfield 2000, 449–450). Typical cases of deferring involve an ordinary speaker and an expert. The deferrer and the expert identify their roles in communication, and they both recognize the existence of an external rule that governs their use of words. Under this description, it is difficult to see how one can accept the idea of ‘unconscious deference’ endorsed by Recanati (2000, 282) – the idea that speakers who do not master a concept unconsciously defer to the linguistic community for its meaning. Furthermore, Woodfield is reluctant to accept that language learning involves semantic deference.² Learning, he holds, is a matter of degree, while the deferential operator, which, according to Recanati, must accompany every use of an imperfectly mastered concept, cannot disappear gradually.³

We find, then, a cluster of problems around the notion of semantic deference. To address those problems, one had better start by identifying the different cases in which we would want to talk of deference. The tendency has been to treat all those cases in terms of semantic deference. This, we argue, is a mistake. Our aim in this paper is to clear up potential confu-

² Willing to take up these challenges, though, Recanati tried to provide a model of concept learning that would incorporate both the idea of gradual improvement and an explanation of deference in terms of the deferential operator (2000, 282–285).

³ Note that Burge explicitly excluded learning from his explananda (1979, 90). A child who memorizes “ $e=mc^2$ ”, he held, does not have beliefs about the theory of relativity.

sion by articulating the differences among several phenomena all of which seem to deserve the label 'deference'.

1.2 The plan

In our taxonomic study, we draw a distinction between linguistic and epistemic deference. Within linguistic deference itself, we distinguish between default deference and deliberate deference. A first approximation to those distinctions may be provided using a single example, the arthritis example, a version of which was given by Burge (1979). Additional examples will be provided as we get to discuss those distinctions in more detail.

Consider a woman who, coming back from the doctor's, tells her partner "I have arthritis." Although it is not the first time she has heard of arthritis, she only has a vague idea of what arthritis is, insufficient for distinguishing arthritis from many other diseases. Thus, she may be unable to differentiate between arthritis, which is a condition of the joints, and myositis, which is a condition of the muscles, and she may even say such things as "I have arthritis in the thigh." Even though the woman's concept of arthritis is poor and, so to speak, indeterminate, i.e. insufficient for fixing the truth conditions of her utterance, this utterance has a determinate truth value, as Burge, Putnam, and externalists in general have successfully argued. This truth value is determined by appealing to the experts, and to the linguistic community more generally, regarding the question of what counts as "arthritis".

This general and truly ubiquitous phenomenon corresponds more or less to what we call default deference. A speaker who defers by default most often does not have the intention to defer.⁴ As a consequence, default deference usually goes unnoticed by speaker and hearer. This contrasts with what we call deliberate deference. A speaker who defers deliberately must intend to do so, and her intention must be recognized by her interlocutors.

In general, deliberate deference involves a language-shift. The speaker intends to use an expression in the way in which it is used in some dialect, sociolect or idiolect. She exploits

⁴ Of course, whenever we engage in communication, we implicitly intend to conform to the rules of language use.

various contextual features to enable her interlocutors to recognize her deferring intention and identify the intended deferree. Consider for instance two doctors who have a patient in common, and suppose that this patient believes that arthritis is an inflammatory condition of the muscles, and keeps saying to her doctors things like “I’ve been suffering so much from the arthritis in my left thigh.” Now suppose that one of the doctors has pain in his calves, and, making it clear to the other doctor that he is alluding to their common patient, he says “My calves really hurt. It must be arthritis.” Here, the doctor intends to use the word “arthritis” in the way the patient does, that is, for muscle inflammation. He makes use of the context to elicit a language-shift to the patient’s idiolect when it comes to interpreting the word “arthritis”. Deliberate deferring, in sum, is an intentional act in which the dialect deferred to must be made salient by the speaker and identified by the interpreter.

An important notion that is related to linguistic deference but should not be confused with it is that of epistemic deference. Let us go back to the lady who, coming home from the doctor’s, tells her partner “I have arthritis.” To establish the truth value of her assertion, we need to determine with respect to which language (dialect, sociolect, idiolect) her words must be interpreted. In English, the meaning of “arthritis” is established in connection with the common body of medical knowledge. So, even if the lady has picked up the word “arthritis” from her doctor, it is not quite right to say that she is deferring to him for her use of the word. For imagine that the doctor himself is mistaken on the question of what arthritis is, and believes that it is a condition of the muscles. Then, if the lady says “I have arthritis in the thigh,” with no overt intention to defer precisely to her doctor, her utterance is false, given that arthritis is a disease of the joints. Now, even though the lady defers by default to the English linguistic community, not to the doctor, there is a sense in which she does defer to the doctor. But, rather than to his linguistic competence, she defers to his judgment (his diagnosis) that she has arthritis. This is what we call epistemic deference.

Below, we present more thoroughly the distinction between default and deliberate deference, and we argue that both phenomena are distinct from epistemic deference. We also discuss in more detail the phenomenon of imperfect mastery, already mentioned above. As we will show, partial understanding of a

concept implies neither that a speaker using the word associated with the concept will intend to defer to others for the meaning of that word, nor that a speaker using the concept in making a claim will base this claim on someone else's judgment. Thus, even though they are often not differentiated, default deference, deliberate deference, epistemic deference, and imperfect mastery will be shown to be distinct phenomena.

2 Linguistic deference: default deference vs. deliberate deference

Default deference is involved in every communicative act. When interpreting and evaluating an utterance, we must take into account a language parameter (which is typically the language of a larger linguistic community, like English, though it can also be a dialect, sociolect or idiolect), and this language parameter is contextually given a default value. To designate this default value, we will use the term "source language". Default deference takes place whether or not we seek to defer. Deliberate deference, on the other hand, is something done intentionally by the speaker. The speaker targets a particular value for the language parameter and exploits the context to help the interpreter identify this value. In this section we will illustrate this distinction with a number of examples and provide a more complex theoretical panorama. We will show that, even if a speaker typically defers by default to the linguistic community, she can also defer by default to a sociolect or idiolect. Similarly, in deliberate deference, even though a speaker typically intends to defer to a certain sociolect or idiolect, she can also deliberately defer to the linguistic community.

There are a few distinctions with which the distinction between default and deliberate deference might be confused, so let us forestall those possible confusions before proceeding. Deliberate deference is intentional and therefore conscious: a speaker who is deferring deliberately must be aware of what she is doing. But this does not make the default/deliberate distinction collapse into the self-conscious/unconscious distinction,⁵ for in the case of default deference, too, the speaker

⁵ We are thinking of Recanati (2000, 281ff) here, even though Recanati himself opposes self-conscious deference to imperfect mastery, and does not

may be perfectly aware of the fact that she is deferring by default. A second possible confusion consist in seeing default deference as semantic and deliberate deference as merely pragmatic: in the default case, the truth value of the utterance containing the deferential expression would depend on the source language provided by the context, whereas in the deliberate case, the speaker would merely convey her intention to use an expression in the way in which it is used by the deferee, without this impinging on the actual truth value of the utterance. On our view, however, both default and deliberate deference affect the truth values of utterances.

2.1 Default Deference

When the lady comes home from the doctor's and tells her partner "I have arthritis in the thigh", our intuitions are clear that the truth conditions of the utterance involve arthritis, not any other medical condition. As the lady cannot have arthritis in her thigh, she is saying something false. People often say false things when they use words they do not completely understand. This is the widespread phenomenon that Brian Loar called "falsity-due-to-misunderstanding" (Loar 1990).

Cases like these were used by externalists to show that, if the sense of a term were identified with the set of descriptions available to a competent language user, then this sense could not determine the term's semantic value. Falsity-due-to-misunderstanding is possible only because the terms we use in our utterances make a semantic contribution that is fixed by linguistic conventions that reflect the community's knowledge of the way the world is. In externalist frameworks, this idea is often grounded in a theory of the social division of linguistic labor. In every linguistic community, there are special groups of language users, the experts, who are entrusted with an important task: determining the semantic value of the terms of the language. Average members of the linguistic community defer to these experts whenever they have to determine the truth conditions of utterances like "I have arthritis in the thigh". The words used by the lady in the context described above acquire their semantic value through these experts. As arthritis is a disease of the joints and cannot affect the mus-

speak of unconscious deference.

cles, the proposition expressed by her utterance is false. Our intuitions about the truth conditions of this proposition are justified by her deferential use of the term “arthritis”. Deference bridges the gap between the ‘arthritic’ lady’s incomplete understanding and the way the world happens to be.

The process described above is what we call default deference. As suggested earlier, whenever an utterance is produced for the purpose of communication, the participants in the communicative exchange have to settle on a language with respect to which interpretation can be carried out, i.e. the source language. In cases of imperfect mastery, it is the experts who determine to which thing or event a given expression applies correctly.

If we resort to Recanati’s deferential operator⁶ to analyze what the arthritic lady says in the context above, and use it in the manner suggested by Recanati, we get the following representation:

(1) I have R_{doctor} (arthritis) in the thigh,

where “ R_{doctor} (arthritis)” is the complex expression that results from the application of the deferential operator to the term “arthritis”. The semantic value of this complex expression is arthritis (the actual disease). (1) is false because arthritis is a disease of the joints and there are no joints in the thigh. The subscript specifies who is being deferred to, in this case the doctor whom the lady visited. As we shall see, this kind of analysis does not entirely do justice to our intuitions about the truth conditions of deferential utterances.

2.1.1 Deference by default is not always deference to the “experts at hand”

Let us imagine that the doctor the lady consulted is not a real doctor, but some madman in a white coat who had just es-

⁶ Here is how Recanati initially characterizes the deferential operator: “The deferential operator is translinguistic. It belongs to a certain language, say L, but applies to any symbol s whether or not s also belongs to L. The result of applying the deferential operator to a symbol s, whether or not that symbol exists in L, is a well-formed expression of L, which I write as ‘ $R_x[s]$ ’, where x is a user of the symbol s [...]” (Recanati 2000, 272). Let us note that we will be using the variable x either for users of a language or for the language itself.

caped from a psychiatric ward. This bogus doctor thinks that arthritis is nothing but a bad hangover. Moreover, he has got it into his mind to pay no attention to his lady patient's symptoms and to tell her that she has arthritis. The lady comes home and reports to her partner that she has arthritis. What she says could conceivably be represented with the help of the deferential operator:

(2) I have $R_{\text{bogus doctor}}$ (arthritis).

But are the truth conditions of the lady's utterance correctly captured by (2)? We do not think so. What the lady says is true if and only if she has in fact arthritis. The semantic contribution of the term "arthritis" to the proposition expressed by her utterance is not a bad hangover, as (2) states, but arthritis. Perhaps the lady had one too many glasses of vodka the night before and was suffering from a bad hangover on the day that she uttered (2), but that would not make the proposition expressed by her utterance any truer. Only arthritis can make that proposition true.

What is happening here? The lady is deferring epistemically to the bogus doctor, since she trusts his diagnosis without further questioning. But she is not deferring to the bogus doctor for the meaning of the term "arthritis". Instead, she is deferring by default to the norms of the linguistic community. She is not deferring by default to the first expert at hand, but to whoever really knows the meaning of "arthritis". Only this ideal expert can satisfy both the externalist claim that meanings are in the world and our intuitions about the truth conditions of her utterance in this context.

2.1.2 Deference by default is not always deference to the linguistic community

Pedro and María are watching the race walking competition in Beijing 2008. Pedro has not seen a race walking event in his entire life, but María, who knows a thing or two about the rules, has just spelt out to him the difference between walking and running in this Olympic sport. Some time after the start, the following dialogue takes place:

- (3) (a) Pedro: “Hey, the second guy is walking so fast he’s gonna catch up with the one in the lead!”
- (b) María: “Actually, he’s running... I’d say he’s gonna be disqualified”.
- (c) Pedro: “Oh, yes, you’re right, he had both feet off the ground for a fraction of a second.”

If Pedro and María were not in a race walking context, their judgments would probably be different from those expressed in the above conversation. For instance, it is not unreasonable to assume that neither Pedro nor María would distinguish between the first and second contestant, so similar is the way they are moving. It is quite possible that they would judge both to be running rather than walking. Yet, these are realistic assumptions only if we take Pedro and María to be using the verbs “walk” and “run” in their ordinary sense. And our claim is precisely that they are not.

We think that, in the context at hand, the source language is not the common language but the particular sociolect of the race walking community. This community has its own experts, namely IAAF judges. These experts define walking in their rule 230 as “a progression of steps so taken that the walker makes contact with the ground, so that no visible (to the human eye) loss of contact occurs”. Rule 230 is the convention that determines the correct application of the term “walk” in this context. The judgments expressed in (3), and the distinctions underlying them, only make sense with respect to such conventions.

Various elements contribute to making the race walkers’ sociolect the source language here. Pedro and María are watching an Olympic race walking event, a sport whose rules they are now familiar with. They have been talking for some time about the technical interpretation of terms like “walking” and “running”. The meaningfulness of their conversation, including their initial disagreement, is further evidence that they are not speaking everyday English. We can therefore conclude that deference by default can select a source language whose conventions differ from those of the language community as a whole.

2.1.3 Deference by default to a particular ‘local dialect’

Imagine twin sisters, Natalya and Olga, who have been brought up in a very isolated area. Their parents use Standard English, except in one respect: they have a peculiar sense of humor, and thought it would be fun to always use “apple” for “pear” (and conversely) in their daughters’ presence. This is a reclusive family and, by the age of six, the sisters have hardly had any contact with anyone outside the family. On their first day at school, the two six-year-olds share the meal their parents have prepared for them, including some fruit. Looking enviously at her sister’s bigger pear, Natalya says to Olga:

(4) Hey, that’s a huge apple!

Any speaker of Standard English would say that the fruit is a pear and would therefore judge Natalya to have uttered a false proposition. But the thing is, in this case, that there is no speaker of Standard English involved in the situation. Both the speaker and her addressee are using the local dialect that their parents have taught them. The whole of their linguistic community actually amounts to themselves and their parents (in their playful mood). The sisters are not even aware that there is a wider linguistic community whose norms may differ from what they have learnt from their parents. Thus, when Natalya or Olga use “apple” and “pear”, they defer by default to their parents’ invented dialect, not to the norms of a language community of which, strictly speaking, they are not part. Using the deferential operator, one could represent the proposition expressed by Natalya and understood by Olga as:

(4’) The object Natalya is pointing at is a huge R_{parents} (apple),⁷

which is the same as:

(4’’) The object Natalya is pointing at is a huge pear.

⁷ We are not, at this stage, claiming that the deferential operator is suitable for analyzing instances of default deference. Here, it is simply a convenient means for representing the manner in which “apple” is to be interpreted. See section 2.1.4 for a discussion of some problems raised by the application of the deferential operator in default cases.

We are aware that this analysis is not self-evidently the right one. All the same, we believe it to be plausible: Natalya and Olga have always deferred by default to their parents, who were the purveyors of the linguistic norm in their environment.⁸ As long as their linguistic community does not extend beyond their parents and each other, they could not defer to anyone else than their parents. This situation will change if their conversation is overheard by someone who knows for a fact that (according to the conventions of Standard English) the fruit they were talking about is a pear, not an apple, and who feels it is her duty to set the record straight for the kids. If the children accept that they are dealing with someone who is more trustworthy than their parents, they will probably change their minds about apples and pears. In our framework this change of mind would be explained as follows: the sisters will have realised that they belong to a wider language community, and that in the community there are experts who are more reliable (more knowledgeable) than their parents as to what this or that object or event ought to be called. If that is what takes place, then the source language of their utterances will shift from their parents' local dialect to Standard English. Now, next time Natalya says to Olga that the fruit she is pointing at is an apple, her utterance will no longer be true if the fruit is indeed a pear. Thus, we see that, other things being equal, the truth conditions of an utterance of (4) are affected by a change in the source language.

2.1.4 Woodfield vs. Recanati on deference by default

Some of the issues we have raised here shed light on certain aspects of the debate between Andrew Woodfield and François Recanati. Woodfield's opposition to an unrestricted application of the deferential operator is partly rooted in his conviction that in cases of partial understanding the expert is not the final source of normativity the deferrer is seeking for:

⁸ Had the conventions of their parents' language been the same as those of the common language, Natalya and Olga would *eo ipso* have deferred by default to the linguistic community as a whole. But the point here is precisely that (a few of) the conventions set by the parents clash with those of the whole linguistic community.

“Both parties [deferrer and expert] take for granted that there are norms which determine the proper meaning of the word, norms to which they both owe allegiance. D [the deferrer] defers to E [the expert] on a particular issue because D takes E to be a good guide, given the meaning that the word already has. D does not take E to be the giver of meaning. No fact about E constitutes the word’s meaning what it does. D knows that experts are fallible. D regards E’s judgement as good evidence that the word means such and such, but D does not suppose that E makes it the case that the word means such and such”. (Woodfield 2000, 450)

The example of the bogus doctor in 2.1.1 gives support to this general intuition. Deference does not always convert some speaker into a ‘giver’ of meaning. In our example, even if the lady heard the word “arthritis” for the first time from the bogus doctor, that would not mean that the bogus doctor could impose his peculiar use of “arthritis” on the meaning of the lady’s utterance. The intuitions about the truth conditions of the proposition expressed in that case are that the semantic contribution of the word “arthritis” is the disease arthritis, and not a bad hangover.

To make his point, Woodfield (Woodfield 2000, 448) resorts to the following example. Alf is a boy who has been told by his teacher that Cicero’s prose is full of synecdoches. The boy picked up the word “synecdoche” from his schoolteacher, unaware that the latter systematically called “synecdoches” what are actually metonymies. Alf meets L, an expert who knows what a synecdoche is, and the following conversation takes place:

- (5) (a) Alf: “Cicero’s prose is full of synecdoches.”
 (b) L: “No it is not. It’s true that his prose is full of figures of speech. But very few of them are synecdoches.”
 (c) Alf: “I accept what you say. Cicero’s prose is not full of synecdoches.”

According to Recanati, Alf’s utterances should be analyzed in the following way:

- (5) (a') Cicero's prose is full of R_{teacher} (synecdoches)
 (c') Cicero's prose is full of R_L (synecdoches)

On this view, what " R_{teacher} (synecdoches)" contributes to the proposition expressed by Alf in (5a') is the content the teacher attributes to the word "synecdoches", that is, metonymies. In (5c'), however, the semantic contribution of " R_L (synecdoches)" is synecdoches. It would seem then that Alf and L are 'talking at cross-purposes'. What Alf says in (5c) does not deny what he said in (5a). This, for Woodfield, is an unacceptable situation.

Within our framework, this problem does not arise. When Alf says in (5a) that Cicero's prose is full of "synecdoches", he is not using his schoolteacher as a sense-giver; he is deferring by default to the linguistic community through his teacher, whom he takes to be a reliable expert. As in the case of the bogus doctor, deference by default is not always deference to the first expert at hand. Alf is deferring to the linguistic community, and thus the semantic contribution of the word "synecdoches" as used in (5a) is synecdoches, not metonymies. When Alf is corrected by L, he learns something about Cicero's prose. His utterance in (5c) is the negation of (5a), because the semantic contribution of the term "synecdoches" is, in both cases, synecdoches. In (5a) he was deferring epistemically to his teacher, since he was taking for granted what the teacher had told him. Now that he has found a more reliable source of knowledge, he decides to defer epistemically to this new source, namely L, and consequently reconsiders his first statement that Cicero's prose is full of synecdoches.⁹

Our discussion of (5) shows that there is a difficulty with the application of the deferential operator to instances of default deference. The question is whether the deferential operator can be used in a proper representation of the truth conditions of utterances like these while, at the same time, preserving the explanatory power that the device has for cases of deliberate deference.

This is a problem that Woodfield has successfully detected.

⁹ We are ignoring the possibility that Alf might be deferring deliberately to his teacher, in which case the proposition expressed by his first utterance would contain metonymies, not synecdoches. This is not the way the example was originally framed by Woodfield.

However, his own characterization of deference faces real difficulties. Though it is adequate for what we have called default deference to the linguistic community, it would have a hard time accounting for cases in which the speaker defers to a certain sociolect or local dialect, as illustrated by the examples of the race walkers and the misled sisters. Given that Woodfield recognizes only deference to the linguistic community, it is not clear how he could deal with these examples, where it is plausible to assume that the IAAF judges and Natalya and Olga's parents play a central role in fixing the meaning of certain terms. Furthermore, Woodfield would have some difficulty accounting for deliberate deference too, since typically, as we shall see, one defers deliberately to a certain idiolect or sociolect, rather than to the whole linguistic community.

2.2 Deliberate deference

Imagine that Tineke and Jan know about the bogus doctor who takes arthritis to be a bad hangover. And they like the story. Last night they partied especially hard and had a lot to drink. In the morning, they wake up and Tineke says to Jan:

- (6) Jan, I have this bad case of arthritis. Would you close the curtains and hand me some aspirin?

Tineke's head is aching badly and she is feeling sick. But she is making playful use of the bogus doctor's misapplication of the term "arthritis" to say that she has a hangover. She knows she can rely on certain contextual features to make manifest the language with respect to which the term "arthritis" is to be interpreted. Tineke goes even further than that: she engineers a language-shift to a target language different from the source language set by default, namely a shift from Standard English to the bogus doctor's idiolect. This, she can afford to do because she can rely on certain features of the context, notably the fact that Jan and herself had a laugh about the story of the bogus doctor and had a lot to drink the previous night. But Tineke could also have made her meaning clear by uttering "I have a bad arthritis", thus using, anomalously, the mass noun "arthritis" as if it were a countable noun. Contextual features of the above kind pretty much ensure that (6) is going to be

understood by Jan as expressing the following proposition:

(6') Tineke has a bad case of $R_{\text{bogus doctor}}$ (arthritis).

Unlike what we observed in cases of default deference, the application of the deferential operator does not raise any issues here. Thanks to it, we can show how Tineke managed to express the proposition that she had a bad hangover, even though she uttered the word “arthritis”. (6) is a paradigmatic case of deliberate deference. It presents all the characteristics of instances of ‘self-conscious’ linguistic deference mentioned in the literature: (i) the speaker chooses to defer for the interpretation of some of her words; (ii) she defers to someone’s idiolect; (iii) deference takes the form of a language-shift that results from the exploitation of certain contextual features. It is to the needs of these examples that Recanati’s deferential operator is tailored.

As we shall see, the above characteristics are not exhibited by all cases of deliberate deference. The following sections are devoted to a scrutiny of non-paradigmatic instances of deliberate deference. In sections 2.1.2 and 2.1.3, we established that the status of the source language (common language vs. sociolect and local dialect) was not constitutive of default deference: a speaker can defer by default not just to the whole linguistic community, but also to a sociolect (the race walking example) and even to a very local dialect (the misled sisters example). Similarly, we will see in the next section that characteristic (ii) does not apply systematically: deliberate deference does not have to involve a shift to someone’s peculiar idiolect; in some cases, the target language is a sociolect or even a common language like English. This means that the distinction between default and deliberate deference cannot be a matter of the sort of language to which speaker and hearer defer. In section 2.2.2, we will show that characteristic (iii) is not a necessary condition for deliberate deference either. We therefore propose an account of deliberate deference that does not appeal to language-shifts in the strict sense.

2.2.1 Deliberate deference to the linguistic community
Suppose that an interdisciplinary wild bunch are working fran-

tically on a taxonomy of linguistic deference. For several hours now they have been discussing similarities and differences between certain examples of default deference and borderline cases of deliberate deference. All the participants, A, B, C and D, agree on a common characterization for these terms and are now trying to tie up the remaining loose ends. The debate seems never-ending. At a critical moment, realising that lunchtime is almost over, the most obstinate, A, tells the others:

(7) All right, let's say that, in deference to you, I'll accept your argument.

We assume that the source language of their discussions is a local dialect that conforms to the definitions on which they had previously agreed. But, if A's utterance is understood by B, C and D, they will not think that A is deferring linguistically or epistemically to any of them, but rather that she is accepting their argument out of respect for them. However, respect is not what the word "deference" would mean in the source language of this context: it is a meaning it has in a different language, namely Standard English.

In (7), the speaker again exploits contextual features in order to make it manifest that she means to shift out of the source language (the deferentialists' dialect) and into a target language that is Standard English. The co-text plays a central role: in its source-language technical sense, the noun "deference" does not collocate with "in ____ to you". This alone should be enough to induce recognition of the shift. All in all, this example shows that deliberate deference does not necessarily rest on language-shifts to an idiolect or a sociolect. One can defer deliberately to the linguistic community.

2.2.1.1 Deference, polysemy and Humpty-Dumpty

This, at any rate, is the conclusion if our analysis is the right one. Yet, we are aware of another possible account for (7): it could be said that, in uttering (7), A simply exploits the polysemy of the word "deference". In other words, where our analysis posits a language-shift from a technical dialect into Standard English, others might see no shift at all. Their argument, then, would be that the technical dialect of the interdisci-

plinary team is nothing but an extension of the standard language. In this extension, the ordinary senses of “deference” (respect and compliance with another’s judgment) are inhibited, while a technical sense is highly activated. On this view, all A does in uttering (7) is reactivate an ordinary sense of “deference”.

We have some sympathy for this analysis. However, we think that its implications are not so straightforward as they look. First, notice that if polysemy is involved in (7), then it is polysemy of a special kind, for the sense that “deference” has in the source language (the technical local dialect) is not (yet) one that is recorded in the lexicon of the target language (Standard English). The problem here is that the deferentialists’ work results in “deference” acquiring a new meaning. Neologisms and meaning-creations always originate in the margins of the common language. Sometimes they catch on, sometimes they don’t. But, if they do, it is always because some aspects of the language spoken by a small group become incorporated into the common core. Until that happens, those aspects cannot be said to belong to the common language. Actually, as some lexicographers have shown (e.g. Rey-Debove 1978, 283–286), new words and new lexical meanings, when they occur in utterances of the common language, are often set off by quote marks or special prosody, indicating that they still feel like words in another language. Our analysis in terms of deliberate deference provides an explanation for the diachronic process by which lexical creations may become part of the common core. In the case of new meanings, this will lead to increased polysemy, but only after the process of extension of the common language has been completed.

Our analysis is less susceptible to accusations of Humpty-Dumptyism than an account strictly in terms of polysemy. On our view of deliberate deference, a speaker does not decree that this or that expression is to be ascribed a new meaning. Rather, she uses expressions which have already acquired a meaning in a given language (be that a common language, sociolect or idiolect). The only decision the speaker makes is to exploit contextual features in order to induce the appropriate language-shift. That is not Humpty-Dumptyism. By contrast, those who reject the deferential account and argue that examples like (7) exhibit plain polysemy can be suspected of Humpty-Dumptyism. On their view, a single language underlies the

interpretation of (7), namely an extension of the common language. This extension includes a new meaning of an already existing term. It seems then that, merely as a result of their theoretical debates, the deferentialists have succeeded in creating a new meaning for “deference”. This means that they have acted pretty much like Humpty–Dumpty in the Alice story.

2.2.1.2 Two more examples

We have shown how our account could accommodate the intuition that polysemy is somehow involved in (7). However, we believe that there are examples similar to (7) for which a polysemy-based account is not even a likely contender. We present two such cases below. The first illustrates deliberate deference to the linguistic community, while the second shows that a speaker can deliberately defer to another common language.

Imagine a guru who, though using the spelling, the grammar and large chunks of the English lexicon, nevertheless chooses to redefine a whole class of key terms (say, “life”, “love”, “devotion”, etc.) in such a way that the ordinary senses of these terms no longer have currency in the language of the guru’s community. One can hardly say here that the guru’s language is a mere extension of Standard English. Now imagine that the guru is preaching to his flock and that his sermon is broadcast on his own satellite TV channel. For a while, he talks directly to his live audience. At one point, however, he looks straight at the camera and, addressing ‘the rest of the world’, says things like:

- (8) You may experience ‘love’ and ‘devotion’ in your hearts, but these are just debased forms of true love and true devotion.

Our suggestion is that, at least for the interpretation of “‘love’” and “‘devotion’”, the guru shifts into Standard English. This, we indicate by means of scare quotes, to reflect the fact that the language–shift engineered by the guru is a deliberate one.

Examples (7) and (8) belong with a class of utterances which display an intrasentential shift into another common language, as in:

- (9) Barthes described the book as “un choc historique” and “un repère nouveau et un départ pour l’écriture”. (Times Literary Supplement, 03/05/02 : 9)

The shift here may be for the sake of accuracy in quoting, or for local color, or meant as a display of one’s linguistic skills. However that may be, this is a deliberate language–shift into a common language: French. Although we acknowledge that (9) is different from the previous two examples, it provides further evidence suggesting that deliberate deference is not systematically to idiolects and sociolects.¹⁰

2.2.2 Deliberate deference without language–shift?

We now consider a class of utterances that seem to fall under the same category as the previous ones. Yet, they turn out to lack one important property exhibited by the various examples of deliberate deference studied so far.

Let us assume that Kate, who has no training in law, is attending a trial. Both the judge and the defendant’s counsel use terms of art with which she is not familiar. For instance, it is not obvious to her whether the defendant committed a felony, an offence or a misdemeanor.

During a break, while talking about the proceedings with other members of the audience, Kate is trying to determine the sort of crime that the defendant is guilty of. In so doing she says things like:

- (10) I don’t think what he did was a felony. I’d say it was a misdemeanor.

Since she realises that her understanding of these terms is at best sketchy, she often supplements her utterances with a metalinguistic comment, or articulates them with a special intonation pattern, of the sort that can be rendered by means of scare quotes:

- (10’) I don’t think what he did was a felony, as the judge put it. I’d say it was a misdemeanor, if I understand the

¹⁰ There are plenty of examples like (9), and they are usually taken to be related to quotation. (See issue 17 of *The Belgian Journal of Linguistics* for various discussions.)

lawyer's distinction.

(10'') I don't think what he did was a 'felony'. I'd say it was a 'misdemeanor'.

These comments and extra markers indicate that we are not dealing with instances of default deference. So, are we dealing with deliberate deference, and, if so, to whom? In an externalist framework such as ours, it is generally accepted that the meaning of legal terms is fixed by members of the legal profession for the whole of the linguistic community. There should therefore be no difference between the meaning that the judge ascribes to "felony" and "misdemeanor" and the meaning that these terms have in the lexicon of English. And if there were a difference, Kate, as a non-expert member of the audience, would probably choose to trust the norms of the linguistic community (as fixed by the body of experts alluded to above). This suggests that, when Kate utters (10), (10') or (10''), she is not (just) deferring to this judge or lawyer, or even to the legal profession, but to the norms of the linguistic community.

It is tempting to conclude that examples (10)–(10'') are a further illustration of deliberate deference to the linguistic community. But, as hinted above, these examples lack one significant feature exhibited by the other cases: they involve no language-shift. In (10)–(10''), the language with respect to which terms such as "felony" and "misdemeanor" are interpreted is none other than the source language set by default. This entails that, unlike what can be observed in (6) and (7), Kate's deferring turns out to have no impact on the truth conditions of her utterances. Still, there is a major difference between (10)–(10'') and genuine cases of default deference. Kate resorts to metalinguistic comments or special intonation patterns in order to make the language parameter of the context salient. This does not happen in cases of default deference, where the speaker typically has no communicative intention to bring the language of interpretation into the foreground.

Faced with these facts, we believe that the right theoretical choice consists in maintaining that (10)–(10'') involve deliberate deference. Accordingly, we must relax criterion (iii) of paradigmatic instances like (6), so as not to require the pres-

ence of a language–shift in the strict sense.¹¹ We therefore propose the following definition:

S performs an act of deliberate linguistic deference if and only if:

- (a) S produces an utterance u;
- (b) S exploits certain contextual features in order to make salient the linguistic parameter L for the interpretation of u or some segment of u;
- (c) S wants her exploitation of contextual resources to be recognized as part of her communicative intentions by the audience.

Although our definition does not include any requirement for a language–shift, cases of deliberate deference can still be represented by means of the deferential operator. For instance, what happens in (10) can be captured by the following formula:

(10₁) I don't think what he did was a $R_{\text{StandardEnglish}}$ (felony). I'd say it was a $R_{\text{StandardEnglish}}$ (misdemeanor).

The deferential operator indicates that the expressions “felony” and “misdemeanor” are to be interpreted with respect to Standard English. In cases of language–shifts, the only difference is that the value of “x” in “ $R_x(s)$ ” is distinct from the source language. Deliberate deference with a language–shift is nothing more than an important sub–category of deliberate deference.

3 Non–linguistic deference and other related phenomena

In this section, we study the notions of epistemic deference

¹¹ Another option would be to leave the criterion for deliberate deference unaltered. As a result, examples like (10)–(10'') would come under a third category of linguistic deference, intermediate between default and deliberate. In our view, however, these examples have much more in common with deliberate deference; hence we shall not pursue this line of reasoning further.

and imperfect mastery. Epistemic deference should be carefully distinguished from linguistic deference, and our first comments will be focused on justifying this distinction. We then show that the notions of epistemic deference and epistemic evidence, though related, must be kept well apart. Finally, we discuss the phenomenon of imperfect mastery and its relationship to linguistic and epistemic deference.¹²

3.1 Epistemic deference

Deference is an issue of interest not only to linguists and philosophers of language, but also to epistemologists and philosophers of science. It is generally admitted that a lot of the knowledge that we possess is acquired deferentially, by testimony. But deference does not only affect the things we know: it also affects our beliefs, beliefs we are none the less ready to act upon. We receive information from many different sources, and we make choices as to which information to accept and which to reject. Imagine a lady with a rare disease who wants to gather different opinions about her illness before undertaking a medical treatment. Every doctor she meets gives her an opinion, based on evidence and other considerations. If the diagnoses differ, she will have to decide which doctor to trust above all others. But, underlying our beliefs and actions are not just other people's judgments on issues for which there is a fact of the matter. We also defer to others on issues that are largely a matter of personal opinion. Suppose that Takeshi has been told by one friend that Sakura is the best sushi-bar in town, and by another that Mikado is the best. If he wants to take his fiancé(e) for a date, Takeshi will have to decide which friend to trust, whose judgments of taste are more reliable.

In what follows, we will focus on those instances of epistemic deference that underlie assertions, because it is in these cases that epistemic deference may be most easily confused with linguistic deference. We will say that a person who makes an assertion is deferring epistemically when she bases her claim, partly or completely, on someone else's opinion. Typi-

¹² Let it be clear from the outset that our goal in this section is not so much to make a new contribution to the existing literature on epistemic and cognitive issues related to deference, as to clearly distinguish those issues from the ones that arise in relation to linguistic deference.

cally, a person who asserts that she has arthritis is epistemically deferring to the doctor on whose diagnosis she relies. We argue below that this phenomenon is distinct from linguistic deference. Furthermore, it cannot be reduced to the notion of epistemic evidence, even if a certain correlation exists.

3.1.1 Epistemic and linguistic deference

Whereas linguistic deference is involved in fixing the meaning of a term, epistemic deference occurs when a person defers to someone else concerning a particular judgment. Whether a speaker is deferring epistemically or not is independent of whether she is deferring deliberately or by default for the use of the words occurring in her utterance.

To begin with, it is easy to realize that default linguistic deference must be independent of epistemic deference. As we have argued at length, default linguistic deference is a ubiquitous phenomenon. Epistemic deference, on the other hand, occurs when we rest a claim upon other people's opinions. It is not surprising, then, that default linguistic deference can, but need not, co-occur with epistemic deference. For example, suppose that Tim goes to see a doctor who, having examined him, tells him: "You have myositis." Tim does not know what myositis is. He only understands that it is related to the pain he is feeling in his calves. Back at home, he tells his mother:

(11) I have myositis. It is nothing serious. I should just rest for a while.

In saying (11), Tim is deferring by default to the linguistic community concerning the meaning of the term "myositis", and, at the same time, he is deferring epistemically to the doctor, the truth of whose diagnosis he takes for granted. But when Tim tells his mother: "My calves hurt badly", he is certainly not deferring epistemically to the doctor, for he is best placed to judge whether a part of his own body hurts or not. However, Tim will still be deferring by default to the linguistic community concerning the meanings of the words that he is using to report the pain in his calves, such as "calves", "hurt", etc.

Somewhat more interesting are the connections between epistemic deference and deliberate linguistic deference. Let us

approach these through various examples. We have already seen that epistemic deference occurs independently of default linguistic deference. When Tim sincerely asserts that he has myositis, he is deferring epistemically to the doctor, but from a semantic point of view, he is deferring by default to the entire linguistic community. This is even more obvious in his assertion that “it is nothing serious”, which is again epistemically based on the doctor’s judgment, but involves only terms that Tim, a native English speaker, fully masters.

Conversely, deliberate linguistic deference occurs independently of epistemic deference. This is clear from our analysis of example (5) above:

- (5) Jan, I have this bad case of arthritis. Would you close the curtains and hand me some aspirin?

Though Tineke is borrowing the bogus doctor’s deviant definition of “arthritis”, she is not deferring to any medical diagnosis made by that doctor, or even to any opinion that he might have regarding her condition.

The mutual independence of linguistic and epistemic deference can be given a more complex, and more subtle, illustration. Think again of the doctor who intentionally uses the word “arthritis” with the deviant meaning that his patient attributes to it. Suppose that this doctor asked for a specialist’s opinion regarding the symptoms in his calves. Diagnosed with inflammation, which is precisely the condition for which the misguided patient uses the word “arthritis”, he tells the colleague with whom he has that patient in common: “My calves hurt. It is arthritis.” Although he does not defer epistemically for the claim that his calves hurt, the doctor defers epistemically to the specialist for the claim that his condition is “arthritis,” i.e. inflammation of his calf muscles. At the same time, he is deferring deliberately to their patient’s idiolect, for the semantic question of what counts as “arthritis.” In this case, deliberate linguistic deference occurs together with epistemic deference, but with distinct deferees.

In sum, linguistic deference and epistemic deference are distinct and mutually independent phenomena, though they can combine in various ways, as has been amply illustrated in

previous sections.

3.1.2 Epistemic deference and epistemic evidence

The examples given above might suggest that epistemic deference occurs as a direct result of there being insufficient epistemic evidence for making a claim. And it is true that epistemic deference is quite often a matter of the amount of epistemic evidence that one has for making a certain statement. Thus, if you have no independent evidence to assert *p*, but have been told by someone you trust that *p*, you are likely to assert *p*, simply because you rely on that person's judgment. We say in such a case that you are epistemically deferring to that person. On the other hand, when you have the best possible epistemic grounds of your own for asserting *p*, then in asserting *p*, you will probably not want to rely on someone else's judgment.

However, lack or poverty of epistemic grounds are neither a sufficient nor a necessary condition for epistemic deference. Someone who has all the evidence that can be had may still choose to defer epistemically to someone else. Thus consider a doctor who happens to be the greatest expert on arthritis, but lacks self-confidence. It is plausible to say that, when he tells a lady patient "You have arthritis," he is epistemically deferring to his colleagues on the issue of whether that woman's condition is indeed arthritis, even though he has enough of his own evidence for this claim. To indicate that he is doing so, he might say "We believe that what you have is arthritis." Or, imagine that Naïma is a first-rate scientist whose research shows how to achieve cold fusion, but is very shy and insecure. She is doing tests in her lab when a senior researcher, whom she deeply respects and admires, tells her: "You are wasting your time. Believe me, cold fusion is something impossible!" Out of sheer insecurity, she decides to defer to his opinion, even though it directly contradicts a claim that she has excellent evidence for, namely, that cold fusion is possible.

Conversely, there are situations in which people may form and express a firm judgment even on an issue for which they have no good epistemic grounds. Consider a woman whose partner tells her "You have arthritis" just out of some inner conviction. He is not, then, deferring epistemically to anyone.

People do make claims for which they have no good evidence, and which do not reflect other people's opinions. Such claims –people's best guesses, as we might put it– exemplify the case where one lacks epistemic evidence, and yet abstains from deferring epistemically.

3.2 Imperfect mastery

Many philosophers hold that there are concepts, and that concepts are very much like mental files in which information gets stored. Consider the concept that Carmelia has of a certain particular, say François Recanati. Her concept contains three main types of information: perceptual information, e.g. that the particular concerned by this concept is that guy, whom she sees talking right there in front of her, descriptive information, e.g. that he is the author of *Literal Meaning*, and metalinguistic information, e.g. that he is called "François Recanati". Our concepts of universals, too, mostly combine those three types of information. But in many cases, the concepts that we associate with words that we use, like "arthritis", "elm" or "hydrogen", are fairly poor, and the information they contain does not enable us to decide on any given occasion whether the word correctly applies to something we are presented with, or to draw certain inferences that someone more knowledgeable could draw. If the concept that a person associates with a term is poor or, at any rate, not as rich as the concept that experts associate with it, we talk of imperfect mastery. Note that mastery is very much a matter of degree, and that it is not obvious that anyone ever achieves perfect mastery. But to bring the issue home, one might want to know how the phenomenon of imperfect mastery relates to linguistic and to epistemic deference, and ask questions like the following. What information must be present in a concept for one to be able to defer, deliberately or by default, using the associated term? Conversely, could the presence of some information make deferring impossible? Does epistemic deference arise whenever we make assertions using concepts that we do not perfectly master? And will the wealth of information in our concepts prevent us from deferring epistemically?

By way of giving a single answer to these questions, we hold that imperfect mastery is a phenomenon that must be kept separate from linguistic deference and from epistemic

deference. In other words, whatever a person's mastery of the concept associated with some term, whatever the amount and quality of the information contained in the mental file, the following options all remain viable: the person will defer by default when using the term, or she will defer deliberately to some contextually salient dialect, be it or not the dialect from which she picked up the term. Likewise, the person may or may not defer epistemically for assertions that she makes using the term.

3.2.1 Imperfect mastery and linguistic deference

Consider a medical expert whose concept of "arthritis" is as rich and determinate as can be. Does such a person defer linguistically when she uses the term "arthritis"? The intuition is that the meaning of such terms is determined precisely by such experts. So if we say that the expert defers in turn, who could she possibly defer to?

Though there may be a grain of truth in this intuition, our account of default deference does not require the speaker to have the intention to defer, or to know which source language is contextually selected. This means that even our medical expert defers by default to the linguistic community when using "arthritis". It just happens that she is among the experts who ultimately determine the meaning of the term.

We have just shown that "perfect" mastery is compatible with deference by default. Is it also compatible with deliberate deference? Again, the answer is "Yes". To see this, just recall our example of the doctor who suffered from an inflamed calf muscle. This doctor could be assumed to know as much about arthritis as possible. Yet, this did not prevent him to wittily exploit the ignorance of a patient and tell his colleague "It must be arthritis."

In a similar way, imperfect mastery allows both for default and deliberate linguistic deference. A woman who knows virtually nothing about arthritis, except that there is something called "arthritis", can use this word to say true or false things, whether or not she has any intention to defer linguistically at all. Or she may defer deliberately, indicating the source from which she got the word and making it clear that she intends to apply the word to whatever it is that her source applies it to, even though she might have no idea what that is.

3.2.2 Imperfect mastery and epistemic deference

Our level of mastery of a given concept can neither force us to defer epistemically nor prevent us from doing so. Someone who has perfect mastery can still choose to defer epistemically, like the shy scientist who endorses her senior colleague's opinion that cold fusion is impossible. Most often, though, people with excellent mastery of a concept make assertions without deferring to other agents, provided that they have strong enough epistemic grounds for their assertions. It is true, too, that, if we know hardly anything about myositis, we are unlikely to go around making unwarranted claims about it. Thus, if we report that Tim has myositis, we will typically do so because someone whose judgment we trust told us that Tim had myositis, or because we read it in Tim's medical file. In those cases, we defer epistemically. But others with the same level of mastery may make the very same claims without deferring epistemically, e.g. out of some inner conviction, however odd this may seem. In any case, even cautious speakers aware of their poor mastery of a given concept will be ready to make certain assertions about myositis without deferring epistemically. For example, they will confidently assert that myositis is a condition called "myositis", or that they would not like to be diagnosed with myositis, even though they have no idea what that is.

In sum, even though there is probably a correlation between an agent's imperfect mastery of a certain concept and her being inclined to defer epistemically, epistemic deference and imperfect mastery are distinct phenomena, irreducible to one another.

4 Conclusion

In this paper, we have argued for the mutual independence of three related phenomena, namely linguistic deference, epistemic deference and imperfect mastery. One of our initial questions has been what kind of framework could accommodate instances of falsity-due-to-misunderstanding and cases in which a speaker overtly chooses to use an expression the way someone else uses it. Our answer has been that, in both types of cases, an expression or segment of discourse is used

deferentially. This is what grounds the category of linguistic deference. Within this category, we have distinguished two varieties, default deference and deliberate deference, which, together, cover a significant proportion of the examples addressed in the literature.

Default deference has been shown not to be restricted to those cases where the speaker defers to the linguistic community as a whole. We have supplied examples of default deference to a sociolect and even to a local dialect. As for deliberate deference, which is usually understood as deference to an idiolect or local dialect, we have given evidence that it ain't necessarily so. Our examples suggest that speakers sometimes defer deliberately to the norms of the common language. The account we provide differs in one further respect from the picture that emerges from what little literature has been written on the subject. Deliberate deference does not always involve a genuine language-shift.

Concerning the related notions of epistemic deference and imperfect mastery, we have contended that they are distinct from each other and orthogonal to linguistic deference. Someone defers epistemically when they base a claim on someone else's judgment, but this does not entail that they are deferring for the meaning of the words they are using. With respect to imperfect mastery, we have shown that the partial understanding of a concept does not constrain a speaker to defer either linguistically or epistemically.

Providing a taxonomy of the various cases of deference discussed in the literature is like drawing a map of the tip of an iceberg. However accurate the map, it is insufficient. Just as safe navigation requires awareness of what lies under the water surface, any stable theory of deference requires awareness of the philosophical and linguistic issues of titanic proportions that underlie it. In the case at hand, the submerged part of the iceberg comprises issues such as quotation, simulation, echoic uses, irony, polysemy, knowledge acquisition, justification, cognitive architecture and concepts.

References

- Burge, T. (1979). *Individualism and the Mental*. Midwest Studies in Philosophy. Volume IV. Studies in Metaphysics. P. French, T. E. Uehling and H. K. Wettstein. Minneapolis, University of Minnesota

Press

Loar, B. (1990). Personal References. *Information, Semantics and Epistemology*. E. Villanueva. Oxford, Basil Blackwell: 117–133

Recanati, F. (2000). *Oratio Obliqua, Oratio Recta : an essay on metarepresentation*. Cambridge, Mass., MIT Press

Recanati, F. (2002). "Unarticulated Constituents." *Linguistics and Philosophy* 25: 299–345

Rey-Debove, J. (1978). *Le Métalangage: étude linguistique du discours sur le langage*. Paris, Le Robert

Woodfield, A. (2000). "Reference and Deference." *Mind & Language* 15(4): 433–451