

Deferential Concepts: A Response to Woodfield

FRANÇOIS RECANATI

1.

When someone uses an expression deferentially, he relies on others: the word in his mouth expresses the content it has for someone else. Thus, in talking to my teacher I may use the word 'synecdoche', which I only imperfectly master, while deferring to her to provide a proper interpretation for that expression. (For example, I may ask her: 'Is it true that Cicero's prose is full of synecdoches? And what does that mean?')

For such a use to be grounded, the person to whom the speaker defers must herself master the relevant concept, or she must herself be in a position to defer to someone who does, or to someone who defers to someone who does—and so on and so forth. More concisely put:

(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.

This follows from the simple fact that a deferential use is a *parasitic* use, asymmetrically dependent upon other people's use.

Woodfield (this issue) rejects the Groundedness Thesis. He mentions several cases in which one defers to the judgement of others who themselves defer to others without any clear stopping point. But this only shows that different notions of deference are at stake. (Indeed that is one of Woodfield's points: there are several sorts of deference, not all of which have the property that 'deference must end somewhere'.) The only case that matters as far as the Groundedness Thesis is concerned is the case in which the very *content* of what one says or thinks depends upon others.

That there are distinct notions of deference, indeed distinct phenomena with quite different properties, has been emphasized by Keith Donnellan in a paper published some years ago (Donnellan, 1993). Consider the case of 'water'.

Address for correspondence: CREA, École Polytechnique, 1 rue Descartes, 75005 Paris, France.

Email: recanati@poly.polytechnique.fr

We all know what water is—we all possess the concept of water—and are thereby able to apply the word in normal situations. We can tell water from gin and from olive oil. Yet, if transported to Twin-Earth, we might well mistake the transparent liquid found there for water, even though it is not water, but XYZ. We have to rely on chemists to determine the extension of ‘water’ in such situations, even though in normal situations we are able to apply the word correctly. Contrast this with the elm/beech case. Putnam confesses that he does not know what an elm is; he cannot tell elms from beeches. Neither can I. We both have to defer to others not merely to determine the extension of ‘elm’ on Twin-Earth, but also to apply that word in the most normal environments.

Two distinct abilities are clearly involved in those examples. First, there is the ability to single out paradigms in the normal environment—‘local paradigms’, as Donnellan says. There is a close connection between possessing a concept (the concept of water, or the concept of an elm) and having that ability. Second, there is the ability to determine the extension of the word in any possible circumstance, including Twin-Earth. Even if one knows what water is in the ordinary sense, still one may not be in possession of the relevant chemical criterion. At this point one must defer to experts.

So there are two sorts of deference as far as natural kind terms are concerned.¹ In the ‘elm’ case, one who does not possess the concept of an elm has to defer to others to discriminate elms from beeches in one’s normal environment. In the ‘water’ case, one possesses the concept of water but one has still to defer to others to scientifically determine the essence of water; hence the extension of ‘water’ in all possible circumstances (Table 1).

Table 1

	—	Deference1 (<i>concept</i>) →		—	Deference2 (<i>hidden nature</i>) →		Extension
Expression ‘water’ ‘elm’			Paradigms				

It is the first sort of deference that is relevant to the Groundedness Thesis. We defer to experts to tell us what the hidden nature of ‘water’ is, but we don’t defer to them in the way we defer to others when we don’t know how to use a word. In this case we do not possess the concept in the first place: we are not even able to single out local paradigms. Someone in that situation has to rely on others as far as the use of the word is concerned. *That* process of deference must end somewhere. If everybody uses ‘elm’ deferentially, the word will never be grounded and will never denote a particular tree.

¹ The first type of deference can affect *any* word, whether it is a natural kind word or not (i.e. whether or not it is susceptible to the other sort of deference).

2.

Instead of contrasting ‘the concept of water’ on the one hand and scientific knowledge about the underlying nature of water on the other hand, we can draw a distinction between *two concepts of water*: the ordinary, basic concept which we all possess, and the scientific concept which only some of us possess.

These different concepts arguably correspond to different sets of alternatives among which to single out water. The ordinary concept of water enables us to tell water from vodka and the other liquids to be found in our environment; but we need a more refined concept if we are to tell water from XYZ. The set of relevant alternatives is simply broader in the latter case than it is in the former (Stalnaker, 1993, pp. 219ff.). Though attractive, this view wrongly suggests that our ordinary concept of water does not determine extension with respect to counterfactual circumstances in which the liquid which has all the phenomenal and functional properties of water is not water. (Or perhaps it determines the wrong extension: perhaps our ordinary concept of water would make XYZ water, on the suggested view.) But one of Putnam’s points is precisely that our ordinary, basic concept of water does correctly determine extension even with respect to such counterfactual circumstances; though it does not do it by itself, but in conjunction with a context. With respect to our context (Earth), our basic concept of water singles out the paradigms whose underlying nature fixes the extension of the concept in all possible circumstances, including Twin Earth. It follows that *the basic concept has the same extension as the scientific concept*. The only difference between the two concepts is this: in contrast to the scientific concept, the basic concept is indexical. Water is *that* stuff in *our* lakes and rivers which quenches thirst etc. It is because of this indexical component that the context comes into play in fixing the extension, via the underlying nature of the contextually accessible paradigms.

What happens when we don’t even possess the basic concept? When we only know that elms are (deciduous) trees? There are two options here. Either we say that in such a case we lack a concept to associate with the word; or we posit a third type of concept beside the basic concept and the scientific concept.

Let us briefly consider an argument in support of the first option. It’s an essential property of concepts that they have, or determine, an extension. Now when I know almost nothing of elms, there is a sense in which what I know is not sufficient to determine the extension of ‘elm’. Hence it seems that I do not really possess a concept of elm. That argument is unconvincing, for a reason which should be clear by now. When I possess the basic concept of water, ‘what I know’ is not sufficient to fix the extension either; for the extension depends upon the underlying nature of the paradigms, which may be unknown to me. Yet the basic concept of water gives me contextual access to the paradigms, the underlying nature of which determines the extension. So the basic concept *has* a determinate extension (in context)—an extension which depends not merely upon what is in the head of the thinker, but also

upon the environment. Similarly, consider what someone knows who is ignorant of elms but has acquired the word 'elm'. He knows that elms are (deciduous) trees called 'elms'. This, too, is a concept with a determinate extension. In the context of a given linguistic community, there is a type of tree which is called 'elm' by the more knowledgeable people. The speaker who possesses only the concept in question does not know which type of tree that is, but this no more prevents his concept from having a definite extension than the speaker's scientific ignorance prevents his basic concept of water from having a determinate extension.

I conclude that there are three distinct concepts which may be associated with a word like 'elm'. There is the basic concept, the scientific concept, and a third type of concept which I call a deferential concept. The basic concept and the scientific concept have the same possible-worlds extension. What about the deferential concept? Once again there are two options. We may construe deferential concepts as having a content similar to that of a metalinguistic description ('the tree called an *elm*'). If we do so, the possible-worlds extension of the deferential concept will be different from the possible-worlds extension of either the basic concept or the scientific concept. For there certainly are possible worlds in which, for example, beeches are called 'elms'. On this construal, what I called the deferential concept of elm will not really be a concept of elm. It will not rigidly refer to elms. Alternatively we may construe deferential concepts as having the same possible-worlds extension as the corresponding basic and scientific concepts. Deferential concepts will thus be treated as indexical, like basic natural kind concepts. In both cases the possible-worlds extension of the concept is relative to the context in which it has been acquired. It is the nature of the local paradigms, or the nature of the trees actually called 'elms' in the context of acquisition, which determines the extension of the concept in *all* possible circumstances.

It is the second option which I want to pursue (section 3). It, and it alone, enables us to maintain that there are three *genuine* 'concepts of elm': the basic concept, the scientific concept, and the deferential concept.

Among those concepts, the basic concept seems to be especially important since it is the concept whose possession is the minimum required for membership in the linguistic community. Possession of the scientific concept is not necessary, while possession of the deferential concept is not sufficient. But that is debatable. According to Putnam, 'the nature of the required minimum level of competence depends heavily upon both the culture and the topic' (Putnam, 1975, p. 249). Water is very important in our daily lives, so one is not considered as understanding the word 'water' unless one possesses the basic concept of water, i.e. unless one is able to apply the word in normal circumstances. To do so one has to be able to recognize water etc. But molybdenum is not as important as water for us, hence having the relevant recognitional abilities with respect to molybdenum is *not* required for membership in the linguistic

community (Putnam, 1975, pp. 247ff.). There is a similar contrast between 'elm' and 'tiger', according to Putnam:

In our culture, speakers are required to know what tigers look like . . . [They] are *required by their linguistic community* to be able to tell tigers from leopards; they are not required to be able to tell elm trees from beech trees. (Putnam, 1975, p. 249)

A purely deferential use of the word 'elm' therefore 'passes muster', i.e. meets the rather low standards set by the linguistic community for having acquired that word. For words like 'water' and 'tiger', the standards are much higher: if one is not able to recognize water or tigers, one's use of 'water' and 'tiger' will not pass muster.²

If Putnam is right, two sub-categories must be distinguished within what I called the first sort of deference. Sometimes one possesses only a deferential concept, but that is fine as far as membership in the linguistic community is concerned. The linguistic community does not require that one possess more than the deferential concept. That is what the thesis of 'division of linguistic labour' amounts to. But when the linguistic community requires of its members possession of a basic concept, it is still possible for someone not to possess that concept and to defer to others. Someone may have incomplete mastery of the basic concept of tiger: he may know simply that tigers are large felines, called 'tigers'. The concept possessed by such a speaker is like the concept of elm possessed by people like Putnam and me: it's a deferential concept. Possession of such a concept is sufficient for understanding the word 'elm', however, but (according to Putnam) it is not sufficient for understanding 'tiger'.

Even though deferential concepts have as determinate an extension as basic or scientific concepts, and are sometimes such that their possession is the only thing required for membership in the linguistic community, still there is a clear sense in which they are not self-sufficient. The major difference between deferential concepts and basic concepts is this: the basic concept has a definite extension even if no one in the linguistic community knows the scientific facts about the underlying nature of the paradigms, or if the scientists to whom ordinary speakers defer in those matters themselves defer to other scientists, and so on indefinitely. The basic concept (like the scientific concept) is a self-standing concept. But the deferential concept has an extension only because there are users in the community who do *not* rely upon others and who possess a basic concept, one which gives them contextual access to the paradigms. The deferential concept therefore is not a self-standing concept; it displays

² Putnam's distinction opens up the possibility that in a literate and scientifically minded society like ours, the standards for a word like 'water' might become so high that knowing the chemical composition might become part and parcel of what it takes to understand the word. At least we may imagine a community in which that is so.

the type of asymmetric dependence upon other concepts which justifies the Groundedness Thesis.

3.

In what follows I will ignore the distinction between imperfect mastery at the individual level and the collectively licensed form of imperfect mastery corresponding to Putnam's division of linguistic labour. I will assume that possession of the basic concept is, in general, required for full understanding of an expression. The cases in which the speaker possesses only the deferential concept will all be considered as instances of imperfect mastery in a generalized sense.

Let us consider a well-known example (due to Burge). A woman goes to the doctor. The doctor, after clinical examination, comes to entertain the thought: 'She has arthritis'. He expresses that thought by telling her: 'You have arthritis'. The woman goes back home and tells her husband: 'I have arthritis'. It seems, at this stage, that the woman has acquired the belief that she has arthritis. But suppose, as Burge did, that the woman has only a very vague, possibly mistaken notion of what arthritis is. Let us go even further: suppose she has no idea what arthritis is—she lacks the concept altogether. Has she really come to believe that she has arthritis? Has she not, rather, come to believe that she has *some ailment called 'arthritis'*? That is Donnellan's suggestion (Donnellan, 1993, p. 167). On this view, even if the woman goes about repeating 'I have arthritis', and that sentence expresses the proposition that she has arthritis, still that is not what she believes. What she believes would be more faithfully expressed by a metalinguistic sentence: 'I have an ailment called (by the doctor) *arthritis*'. It follows that there is a divergence between the content of the utterance, which depends on social factors (viz. the conventions in force in the public language), and the content of the underlying mental representation. The mental representation is metalinguistic while the public representation is not.

Even though there is a grain of truth in that theory, I think it must be rejected. The view I myself favour does not rest on a distinction between the content of the public utterance and the content of the underlying mental representation. Like Burge, I hold that one and the same proposition, namely the proposition that the woman has arthritis, is both the content of the woman's utterance and the content of her belief. (It is also, of course, the content of the doctor's utterance and the content of the doctor's belief.)

At this point a difficulty immediately arises. How could the woman believe she has arthritis? In order to believe that she has arthritis, she must entertain a mental representation whose content is the proposition that she has arthritis. Now such a representation must contain a constituent (a concept) whose content is arthritis, and we have granted that the woman does not possess the concept of arthritis!

It is here that the distinctions made in the previous section can be helpful. It is true that, if the woman is to believe that she has arthritis, she must entertain a mental representation whose content is the proposition that she has arthritis. It is true also that such a representation must contain a concept whose content is arthritis. But ‘a concept whose content is arthritis’ is *a* concept of arthritis, it is not *the* concept of arthritis—that concept which the woman lacks. When we say that the woman ‘does not possess the concept of arthritis’, we do not talk of any old concept whose content is arthritis: we refer to a specific concept endowed not only with a certain content (arthritis) *but also with a certain character*. That is the *basic* concept of arthritis. The woman does not possess that concept, but she possesses *another* concept of arthritis—a deferential concept.

Insofar as basic natural kind concepts are indexical, they have a character and a content. As Donnellan himself puts it, when we understand a word like ‘water’, we know a ‘semantic rule’ which (in our environment) determines a set of local paradigms whose underlying nature ultimately determines the extension of ‘water’. The content of the concept is the property of being water, but its character is a function which maps the environment in which the concept has been acquired onto that property which the local paradigms happen to instantiate. Similarly, I suggested that deferential concepts are indexical and possess both a character and a content. The content of a deferential concept arguably is the same property as that of the corresponding basic concept, but its character is different. The difference between the woman’s belief and the doctor’s (or between the woman’s belief and her utterance) is therefore not a difference at the content level, but a difference in character or mode of presentation.

What is the character of deferential concepts? In my earlier work on this topic I suggested that such concepts are constructed by means of the ‘deferential operator’ (Recanati, 1997). The deferential operator $R_x()$ applies to (the mental representation of) a public symbol σ and yields a syntactically complex representation $R_x(\sigma)$ —a deferential concept—whose character takes us from a context in which reference is made to a competent user x of σ , to a certain content, namely the content which σ has for x , given the character which x attaches to σ . What is special with the concept $R_x(\sigma)$ is that its content is determined ‘deferentially’, via the content which another cognitive agent, somehow given in the context, attaches or would attach to σ in the context of utterance.

The deferential operator is the mental equivalent of quotation marks in written speech. It is metalinguistic in the sense that it involves a mention of the symbol σ and a tacit reference to its use by the cognitive agent x (which can be a community as well as an individual). But that metalinguistic aspect is located in the character of the deferential concept: the content of that concept is the same as the content of the symbol σ when used by x .

Take our ‘arthritis’ case. When the woman who does not know what arthritis is says ‘I have arthritis’, she does not entertain ‘the’ concept of arthritis,

i.e. the basic concept, let alone the scientific concept—as the doctor presumably does when he tells her ‘you have arthritis’. They entertain different mental representations, involving different concepts. In the woman’s belief a deferential concept occurs, namely: $R_{\text{doctor}}(\text{arthritis})$. But the content of that concept is the same as the content of the doctor’s concept of arthritis—indeed, the woman’s deferential concept is parasitic on the doctor’s concept and automatically inherits its referential content, by virtue of the mechanics of the deferential operator. That referential content—arthritis—is thought of metalinguistically as ‘what the doctor calls *arthritis*’, but the woman’s thought is fundamentally about arthritis, not about the word ‘arthritis’. If I am right, the difference between the doctor’s concept of arthritis and the woman’s is similar to that between ‘I’ and ‘you’ in their respective utterances ‘You have arthritis’ and ‘I have arthritis’: ‘I’ and ‘you’ refer to the woman under different modes of presentation. Similarly, the doctor’s concept of arthritis and the woman’s deferential concept $R_{\text{doctor}}(\text{arthritis})$ both refer to arthritis, under different modes of presentation.

4.

Woodfield accepts that there is a deferential operator which works in more or less the way I describe, but not my claim that it is at work in examples like the ‘arthritis’ example. The cases that support my view, according to him, are the cases in which we consciously use a word which we do not understand, in quotation marks as it were. In such cases, I hold that the content of the thought or utterance is the same as it would be if no quotation marks occurred and no deference took place: the metalinguistic component is located at the character level. Woodfield is prepared to accept all this. But my theory explains ‘a rather specialized range of phenomena’, he holds. It was a mistake on my part to extend it to cases of imperfect mastery, like Burge’s ‘arthritis’ example. Woodfield thus rejects my claim that ‘children, language-learners, indeed anyone, when they pick up words that they do not fully understand, normally bind such words inside deferential operators’ (this issue, p. 445).

Not only is there a phenomenological difference between self-conscious deference and imperfect mastery; there is, Woodfield points out, a good theoretical reason for not putting them in the same basket. Imperfect mastery is a matter of degree—one’s mastery of a concept is more or less imperfect. In my original article on the topic, I myself insisted that deferentiality is a matter of degree: there is, I said, a continuum of cases between the deferential use of a symbol which we do not understand and its normal use. In between we find instances of partial mastery—as in Burge’s original example. Now this raises a problem for my account, Woodfield says, because

It seems impossible that there should be a *gradual* process of moving out of quasi-quotes. It’s clearly not a process of bit-by-bit removal (like

taking one's clothes off), nor it is a process of decay (like quotation marks fading away on a page as the ink loses its colour). The learner starts off using mental symbols like R_x ('*synecdoches*') and R_x ('*kachna*') and ends up using completely distinct symbols like *synecdoches* and *duck*.³ Prima facie, there has to be a saltation—a switch of symbol-type—at some point. (This issue, p. 447).

I grant Woodfield both points: first, that there is a difference between self-conscious deference and imperfect mastery—a difference which should be accounted for; second, that the gradual nature of imperfect mastery makes it hard, if not impossible, to account for the transition from imperfect to full mastery in terms of a switch of symbol-type. The problem for my account is that such a switch is precisely what adding or removing the deferential operator brings about.

Faced with those difficulties, we may allow for the following possibility. Whenever we mentally entertain a sentence containing a symbol we do not properly understand, the deferential mechanism operates *as if* we had used the deferential operator, that is, as if we had put that symbol within quotation marks and deferred to some authority for its interpretation. But we don't have to actually use the deferential operator—the deferential interpretation can be provided by default, simply because no direct interpretation for the symbol is available to the subject. On this account the difference between conscious deference and incomplete mastery is syntactic, not semantic. In ordinary cases of incomplete mastery, the deferential shift takes place without being syntactically articulated. Since that is so, the continuum from incomplete to complete mastery no longer raises a problem. No saltation needs to be involved because the difference between normal and deferential use no longer lies at the level of the symbol-type. One and the same symbol-type is tokened in both cases. If that symbol is appropriately connected to some concept in the subject's repertoire, it expresses that concept and conveys its content. If the symbol is not appropriately connected to some concept in the subject's repertoire, the concept that is expressed is that which would be expressed by applying the deferential operator to that symbol. On this account, it is only to be expected that the process of connecting up a symbol with concepts in one's repertoire, hence the transition from deference to full mastery, will be gradual.

Though it is a step in the right direction, the foregoing account is not ultimately satisfactory. It violates a principle which I put forward in my original paper, and which we can call the Interpretation Principle:

³ One of Woodfield's examples involves a non-Czech speaker looking at a menu written in Czech and uttering *For lunch I shall have 'kachna'*. 'Kachna' means *duck* in Czech, Woodfield tells us.

(Interpretation Principle)

If a mental sentence is well-formed, it must possess a definite meaning—a character—even if it falls short of expressing a definite content. (Recanati, 1997, p. 91; quoted by Woodfield, this issue, p. 444)

If we accept this principle, then there is an incoherence in the revised account I have just presented. We are to suppose that the subject entertains a mental sentence in which a symbol σ occurs. Whenever that symbol turns out to be uninterpreted by the subject's own lights, it receives a deferential interpretation by default. This violates the Interpretation Principle; for the so-called mental sentence will not be well-formed in the first place—it will not be a *mental* sentence—if it contains some uninterpreted symbol. Mental sentences must be constituted out of the right material: *conceptual* material. The symbols used in thought must be potential conveyors of content; they must be interpreted at least at the character level. That is what the Interpretation Principle requires. The role of the deferential operator was precisely to guarantee satisfaction of the Interpretation Principle. In the same way in which quotation marks can turn a non-word into a well-formed expression of English, the deferential operator can turn the uninterpreted symbol σ into a complex symbol $R_x(\sigma)$, which has a character and possibly a content.

On the revised account, the uninterpreted symbol σ will acquire a character when the deferential interpretation is provided by default. But this is too late: how will the uninterpreted symbol σ come to occur as a constituent in the subject's thought, unless it is already interpreted? This is a serious worry for anyone who accepts the Interpretation Principle.

5.

I suggest that we revise the revised account so as to satisfy the Interpretation Principle. Let us not say that the deferential interpretation is provided by default when an uninterpreted symbol occurs in thought. According to the Interpretation Principle, no uninterpreted symbol *ever* occurs in thought. Still, we want to capture the fact that sometimes, in our thinking, we use a public word which we do not understand. In line with the Interpretation Principle, we want the word in question to receive a deferential interpretation from the very start; and we do not want this interpretation to affect the identity of the symbol-type, as the use of the deferential operator would do. These are the desiderata. To satisfy them, we must give up the view that words are labels associated with concepts. We must *construe words themselves as concepts*, which we can associate with other concepts (e.g. recognitional concepts). Thus, when we acquire a public word, whose use we do not yet fully master, we automatically acquire a concept.⁴ The concept in question is deferential: its content is

⁴ See Millikan 1997, section 6:

It is . . . possible, indeed it is common, to have a substance concept entirely through the

determined via the users whom we get the word from (or via the community in general). When we use a word we do not understand in our thinking, it is the deferential concept which occurs in our thought—hence the Interpretation Principle is satisfied. Again, the public word, insofar as we use it in thought, *is* the deferential concept, it does not have to be associated with a deferential or any other type of concept. In this account there no longer is a gap between the public word which occurs in thought and the deferential interpretation it receives: the deferential interpretation is a built-in feature of public words *qua* thought constituents.

What happens when (gradually) we come to understand the word in a non-deferential manner—when, for example, we get acquainted with what it applies to and progressively acquire the basic concept? We must not think of this process as the association of the word with a concept—an association which was lacking beforehand. Rather it is the association of *two* concepts: a deferential concept and another type of concept. This is the same sort of process which takes place when we *recognize* an object we have seen before: then a past-oriented demonstrative concept ‘that object [which I saw the other day]’ gets associated with a standard demonstrative concept based on current perception: ‘that object [in front of me]’.⁵ In such a situation typically the two concepts get mixed and give rise to a third concept, with a distinct character. Similarly, when a deferential concept—for example, Putnam’s concept of an elm—gets associated with a non-deferential concept (e.g. the demonstrative concept ‘that type of tree’), and that association stabilizes, a new concept

medium of language, that is, in the absence of any ability to recognize the substance in the flesh. For most of us, that is how we have a concept of Aristotle, of molybdenum, and, say, of African dormice. There, I just handed you a concept of African dormice, in case you had none before. Now you can think of them at night if you want to, wondering what they are like—on the assumption, of course, that you gathered from their name what sorts of questions you might reasonably ask about them . . . In many cases there is not much more to having a substance concept than having a word. To have a word is to have a handle on tracking a substance via manifestations of it produced in a particular language community. Simply grasping the phonemic structure of a language and the rudiments of how to parse it enables one to help oneself to an embryo concept of every substance named in that language.

Similar remarks can be found in Kaplan’s ‘Afterthoughts’:

The notion that a referent can be carried by a name from early past to present suggests that the language itself carries meaning, and thus that we can *acquire* meanings through the instrument of language. This . . . provides the opportunity for an *instrumental* use of language to broaden the realm of what can be expressed and to broaden the horizons of thought itself. [. . .] Contrary to Russell, I think we succeed in thinking about things in the world not only through the mental residue of that which we ourselves experience, but also vicariously, through the symbolic resources that come to us through our language. It is the latter—*vocabulary power*—that gives us our apprehensive advantage over the non-linguistic animals. My dog, being color-blind, cannot entertain the thought that I am wearing a red shirt. But my color-blind colleague can entertain even the thought that Aristotle wore a red shirt. (Kaplan, 1989, p. 604)

⁵ See Evans, 1982, chapter 8, for illuminating remarks on this topic.

results, with a distinct character. How is the merging process to be properly described? I do not know, but I have no doubt that it can be gradual, and that is all that matters for us.

From a strictly semantic point of view, there is no significant difference between my original account and the account we arrive at. When in our thinking we use public words which we do not quite understand, our thoughts have deferential concepts as constituents. The character of these concepts is the same as the character of complex symbols built with the help of the deferential operator (as in self-conscious deference). The character in question is metalinguistic, much as the character of indexicals is metalinguistic. Just as 'I' refers to the person who says 'I', 'arthritis', for the patient, refers to what the doctor calls 'arthritis'. But the content of the thought or utterance is metalinguistic in neither case: when she thinks 'I have arthritis', the patient entertains a thought which is about her (not about the word 'I') and about arthritis (not about the word 'arthritis').

If I am right, deferential uses of words are neither deficient nor deviant. They would be deficient if the speaker merely used an empty word to which he or she associated no concept. But that is not the case: deferential uses of words express full-blown concepts, namely deferential concepts. Deferential uses would be deviant if the content of the expression so used was distinct from its normal content—if, for example, the expressed content was metalinguistic, as Donnellan holds. But I argued that the content expressed by a deferential use of a word is its normal content—the difference is located at the level of character. Still deferential uses are parasitic on non-deferential uses. In the context of a linguistic community in which everybody uses the word deferentially, the character of the deferential concept would determine no content. That is the gist of the Groundedness Thesis, which I maintain.

CREA, *École Polytechnique*
Paris, France

References

- Burge, T. 1979: Individualism and the mental. *Midwest Studies in Philosophy*, 4, 73–121.
- Donnellan, K. 1993: There is a word for that kind of thing: an investigation of two thought experiments. *Philosophical Perspectives*, 7, 155–71.
- Evans, G. 1982: *The Varieties of Reference*, ed. J. McDowell. Oxford: Clarendon Press.
- Kaplan, D. 1989: Afterthoughts. In J. Almog, H. Wettstein and J. Perry (eds), *Themes from Kaplan*. New York: Oxford University Press, 565–614.
- Millikan, R. 1998: A common structure for concepts of individuals, stuffs, and real kinds: more mama, more milk, and more mouse. *Behavioral and Brain Sciences*, 21, 55–65.

- Putnam, H. 1975: The meaning of 'meaning'. In H. Putnam, *Philosophical Papers*, vol. 2. Cambridge University Press, 215–71.
- Recanati, F. 1997: Can we believe what we do not understand? *Mind and Language*, 12, 84–100.
- Stalnaker, R. 1993: Twin Earth revisited. Reprinted in R. Stalnaker, *Context and Content*, 1999. New York: Oxford University Press, 210–21.