

Modes of presentation: Perceptual v. Deferential

François Recanati

► **To cite this version:**

François Recanati. Modes of presentation: Perceptual v. Deferential. Albert Newen, Ulrich Nortmann, Rainer Stuhlmann-Laeisz. Building on Frege: New Essays on Sense, Content, and Concept, CSLI Publications, pp.197-208, 2001, 10. <ijn_00000159>

HAL Id: ijn_00000159

https://jeannicod.ccsd.cnrs.fr/ijn_00000159

Submitted on 30 Sep 2002

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.

Modes of Presentation: Perceptual vs Deferential*

FRANCOIS RECANATI

1 Preliminaries

1.1. Two dimensions of content

The content of a representation (be it a concept, a thought, a word, or an utterance) is what it is a representation of. The content of my concept of tiger is the species *tiger*, the content of my concept of Cicero is Cicero. Similarly, I will assume that the content of my thought, or statement, that Cicero is a famous Roman orator is the state of affairs consisting of Cicero and the property of being a famous Roman orator.

But there is a complication. Suppose I heard of Tully in highschool, without ever realizing that Tully is Cicero. I believe that there are two persons, Cicero and Tully, who both were Roman orators. Consider my belief that Cicero was a Roman orator and my belief that Tully was a Roman orator. Do they have the same content? Yes, if the content of a representation is what it is a representation of. For (i) the only difference between the two beliefs is the fact that the concept of Cicero occurs in one while the concept of Tully occurs in the other, and (ii) both concepts are concepts of the same individual.

The problem with this answer is that there is another notion of content which seems equally legitimate. Arguably, two beliefs have different contents if their input-output connections are different. Suppose I believe that Cicero did, while Tully did not, denounce Catiline. Then, my belief that Cicero was a Roman orator will prompt the

* An ancestor of this paper, corresponding to the first two parts, was presented at the *Languages of the Brain* conference organized in Paris in March 1998 by the Mind/Brain/Behavior Interfaculty Initiative (Harvard University) and the Ipsen Foundation. It will be included in the Proceedings of that conference, to be published by Harvard University Press. The present paper, which I read at the Bonn conference, overlaps also with my 'Deferential Concepts: A Response to Woodfield' (*Mind and Language* vol. 15 n°4, September 2000). I am indebted to Steven Davis for helpful comments on an earlier draft, and to ACI Cognitique for funding the last stage of this research.

inference that a Roman orator denounced Catiline; not so with my belief that Tully was a Roman orator. This difference in inferential potential shows that the two beliefs somehow differ in content.

In the familiar terminology derived from Frege (1892), we can say that there are two dimensions of content: *what* is represented (the reference) and *how* it is represented (the 'mode of presentation'). In the Cicero/Tully case, the reference is the same, but the modes of presentation are different. The reference is what is relevant for evaluating a representation as correct or incorrect. In contrast, the mode of presentation — the way the reference is thought of — is relevant for explaining (and predicting) behaviour, including linguistic and mental behaviour.

1.2. Modes of Presentation as Mental Files

What are modes of presentation? How are they individuated? If concepts are construed as mental files, as I think they should be, it will be tempting to equate the mode of presentation with the information (or perhaps misinformation) contained in the file. Thus my Cicero-file and my Tully-file refer to the same individual, but contain different pieces of (mis)information concerning that individual: in contrast to my Tully-file, my Cicero-file contains something to the effect that he (the individual in question) denounced Catiline.

But this qualitative construal of modes of presentation is not ultimately satisfactory. The criterion for distinctness of modes of presentation, as stated by Stephen Schiffer after Frege, is:

(Frege's Constraint)

If it is possible for a rational subject S to believe at the same time of a given object *a* both that it is F and that it is not F, then there are two distinct modes of presentation of *a*, *m* and *m'*, such that S believes of *a* under *m* that it is F and under *m'* that it is not F.

Now for it to be *possible* to believe without irrationality that *a* is F and that it is not F, it is sufficient to have two distinct mental files for *a*, even if the two files contain the same information (e.g. that *a* was a Roman orator). It follows that modes of presentation are not to be identified with the informational content of the mental file, but rather with the mental file itself.

1.3. Egocentric files

One important type of case, much discussed in the recent literature on modes of presentation (e.g. Perry 1993), concerns indexical concepts like *here*, *now*, *this man*, *myself*, etc. What is it to think of an object under such a mode of presentation? What is it to think of myself as *myself*, and not, say, as *François Récanati*? What is it to think of a place as *here*? To think of a pen as *this pen*? Answer: these modes of presentation are 'egocentric files', that is, special files which serve as repository for information *gained in a certain way* — typically through perception. The 'self' file serves as repository for information gained in the first-person way (e.g. through proprioception); the 'here' file serves as repository for information gained about a place by virtue of occupying that place and being in a position to perceive what is going on there. Demonstrative files ('that man') also depend upon perception. In order to entertain a thought in which one of these modes of presentation occurs, one must be suitably related to the reference: one must be related to it so as to be able to gain information from it in the proper way.¹

1.4. Types and Tokens

The relational constraint I have just mentioned is a constraint on the context in which a certain type of mode of presentation can occur. This suggests the following, overall picture. There are mental symbols (types) which are very much like the words 'I', 'here', etc. They can be appropriately tokened only in a context which satisfies certain constraints (having to do with the subject's sensitivity to perceptual information from the object). When the constraints are satisfied, a reference is determined, and the symbol (token) counts as a genuine mode of presentation, that is, a mode of presentation *of* something. The symbols have therefore two sets of semantic properties: *qua* tokens (if the context is appropriate) they have a reference which they present in a certain way. *Qua* types, however, they have a 'character', a constant meaning which corresponds to the constraint itself and determines the way in which the reference, if any, is presented. That constant meaning can be described as a function from

¹ Of course one can have an *I*-thought (or an *here*-thought, or a *this*-thought) in the absence of any perception. (See Anscombe 1975: 58 for a description of such cases.) This shows that the relevant perceptual relation must be thought of in dispositional terms: thus understood, "the informational connection still obtains even if the subject's senses are not operating" (Evans 1982: 161n). For more on this issue, see Recanati 1993: 120-2.

(appropriate) contexts to the referential contents assumed by the symbols in those contexts.²

2 Deference and Indexicality

2.1. Perception and Communication

All the ways of gaining information associated with egocentric files are perceptual ways of gaining information. But perception is not the only mode of acquisition of information: Communication is another. So the question I want to raise in this paper is this: Are there modes of presentation which depend on communication in the same way in which egocentric files depend on perception? My own answer to that question will be 'yes'. But many theorists would rather make a negative answer, and I will start by presenting an argument they might offer in support of their view.

The argument I have in mind goes roughly like this. There is an *asymmetry* between perception and communication. The latter is parasitic on the former, in the sense that whatever information is transmitted through communication must ultimately have been acquired through perception. The thought which occurs at the beginning of the communicational chain must therefore contain perception-based modes of presentation. What happens at the other end of the communicational chain? The recipient of the communicational act does not acquire the relevant information through perception.³ Yet, arguably, he or she must be able to grasp the information thus transmitted, and that means that the information in question must be information she *might* have acquired through perception (even if, as a matter of fact, she acquired it through communication). In other words, the conceptual repertoire which communication exploits must be the same repertoire which we put to use in perception. This supports the claim that there are no modes of presentation which depend on communication in the same way in which egocentric files depend on perception.

As against this view I will suggest that there is a special, communication-based mode of presentation which plays quite a central role in our mental lives *precisely*

² In what follows I will use 'content' always in the sense of 'referential content' (as opposed to 'character').

³ Or at least, not directly. Hearing something on the radio, or from someone, certainly involves auditory perception, but *what* is perceived in such a case is the utterance which conveys the piece of information, rather than the fact stated by the utterance. To hear that John is presently in Paris is not to perceive that he is, but to be told that he is.

because it enables us to go beyond our private conceptual repertoire. I will therefore reject the claim that the same conceptual repertoire is put to use in perception and communication.⁴ But first, I want to spell out a consequence of the view I am going to oppose.

2.2. Deference as Metarepresentation

A woman goes to the doctor who tells her that she has arthritis. She believes what the doctor tells her. A number of representations are involved in this simple exchange. 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 then comes to believe something that might be expressed by the sentence: 'I have arthritis'.

It seems, at this stage, that the woman has acquired the belief that she has arthritis. But suppose, as Burge did in the famous article in which this example was introduced, 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. Does she really come to believe that she has arthritis? Does she not, rather, come to believe that she has *some ailment called 'arthritis'*? This view seems more in line with the notion that whatever is actually communicated must be within the addressee's ken and could not exceed his or her conceptual repertoire.

On the suggested view, then, it is misleading to say that the woman comes to believe that she has arthritis as a result of the doctor's utterance. She does not actually believe that. Even if she 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*' (Donnellan 1993). 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.

2.3. Another Approach to Deference

In contrast to that metalinguistic approach, the view I am going to put forward does not rest on a distinction between the content of the public utterance and the content of the underlying mental representation. On my view, one and the same proposition, namely

⁴ An explicit rejection of that claim can be found in Sperber 1997.

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!

That difficulty is not as dramatic as it seems. 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 are not talking of any old concept whose content is arthritis: we are referring to a specific concept endowed not only with a certain reference (arthritis) *but also with a certain character*. The woman does not possess that concept, but she possesses another concept of arthritis. The concept of arthritis she possesses is a *deferential* concept, i.e. the sort of concept which people who use a public word without fully understanding it typically entertain. 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.

2.4. The Deferential Operator

Qua types, indexical symbols have a character, a constant meaning which determines both their reference (in context) and the way the reference is presented. That constant meaning is a function from contexts to contents. My hypothesis is that there is, in the mental repertoire, a 'deferential operator' which enables us to construct deferential concepts with a semantics analogous to that of indexical concepts (Recanati 1997).

The deferential operator $\mathbf{R}_x(\)$ applies to (the mental representation of) a public symbol σ and yields a syntactically complex representation $\mathbf{R}_x(\sigma)$ — a deferential concept — which has both a character and a content. The character of $\mathbf{R}_x(\sigma)$ 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 deferential concept $\mathbf{R}_x(\sigma)$ 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, as the doctor 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_{doctor}(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_{doctor}(arthritis)** both refer to arthritis, under different modes of presentation.

2.5. Deference as Social Indexicality

On the view I have sketched, deferential concepts are like indexical concepts, but instead of depending on perception they depend on communication. An indexical type of mode of presentation constrains the context of its tokenings: it demands that the subject who entertains the mode of presentation be directly related to the reference so as to be able to gain information from it in perception. Similarly there is a deferential type of mode of presentation which demands that the subject who entertains a mode of presentation of that type (a deferential concept) be indirectly related to the reference, via a chain of informants, so as to be able to gain information from it in communication.

3 Using Public Words in Thought

3.1. Self-Conscious Deference vs Imperfect Mastery

Andrew Woodfield has put forward the following criticism of my view (Woodfield 2000). He 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 Woodfield, are the cases in which we consciously use a word which we do not understand, in quotation marks as it were. In contrast to Donnellan, who holds that in such cases what is believed is a metalinguistic proposition, 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 accepts 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, and other imperfect understanders of picked up words normally bind such words inside deferential operators" (Woodfield 2000: 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, between full mastery of a concept and total lack of that concept. 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 could 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*.⁵ Prima facie, there has to be a saltation — a switch of symbol-type — at some point. (Woodfield 2000: 447).

I grant Woodfield both points: first, that there is a difference between self-conscious deference and imperfect mastery; second, that the gradual nature of imperfect mastery makes it hard, if not impossible, to account for the transition from imperfect to full

⁵ 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.

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.

3.2. Deference By Default

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.

3.3. The Interpretation Principle

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

It is hard to think of a symbol being mentally entertained without being 'interpreted' in some fashion or other... 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: 91, quoted in Woodfield 2000: 444)

If we accept this principle, as I do, then there is an incoherence in the revised account presented in the last section. 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.

3.4. Words as Concepts

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.

I suggest the following account as a means of satisfying the desiderata. We must give up the Aristotelian view that words are labels associated with concepts. We must construe words themselves as concepts,⁶ which we can associate with other concepts

⁶ I am not suggesting that we should identify concepts with words. Obviously, there are many concepts to which no word corresponds. My point, rather, is that a sub-class

(e.g. recognitional concepts). Thus, when we acquire a public word, whose use we do not yet fully master, we automatically acquire a concept.⁷ The concept in question is deferential: its content is 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

of our concepts is constituted by public words which we imperfectly understand but which we are able to use in our thinking.

⁷ See Millikan 1998, §6:

It is... possible, indeed it is common, to have a substance concept entirely through the 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 nonlinguistic 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: 604)

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? 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]'.⁸ In such a re-identification situation typically the two concepts coalesce, are merged into a single recognitional concept, with a distinct character. (That arguably is a *third* concept, distinct from the first two as a child is from mother and father.) 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 results, with a distinct character. How is the merging process to be properly described? I don't know, but I have no doubt that the merging process can be gradual, and that is all that matters for us.

3.5. Conclusion

What is left of my original account? From a strictly semantic point of view, there is no significant difference between the 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 up with the help of the deferential operator. 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').

References

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

- Anscombe, E. 1975. The First Person. In S. Guttenplan, ed. *Mind and Language*, Oxford: Clarendon Press, pp. 45-65.
- 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.
- Frege, G. 1892. On Sense and Reference. In *Translations from the Philosophical Writings of Gottlob Frege*, by P. Geach & M. Black. 2nd ed., Oxford: Clarendon Press, 1960, pp. 56-78.
- Kaplan, D. 1989. Afterthoughts. In J. Almog, H. Wettstein and J. Perry, eds. *Themes from Kaplan*. New York: Oxford University Press, pp. 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.
- Perry, J. 1993. *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press.
- Recanati, F. 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Recanati, F. 1997. Can We Believe What We Do Not Understand? *Mind and Language* 12: 84-100.
- Schiffer, S. 1978. The Basis of Reference. *Erkenntnis* 13:171-206.
- Sperber, D. 1997. Intuitive and Reflective Beliefs. *Mind and Language* 12: 67-83.
- Woodfield, A. 2000. Reference and Deference. *Mind and Language* 15: 433-451.