



Meaning and Ostension: From Putnam's Semantics to Contextualism

François Recanati

► To cite this version:

François Recanati. Meaning and Ostension: From Putnam's Semantics to Contextualism. 2005.
<ijn_00000599>

HAL Id: ijn_00000599

https://jeannicod.ccsd.cnrs.fr/ijn_00000599

Submitted on 20 Mar 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.

**Meaning and Ostension :
From Putnam's semantics to Contextualism**

François Recanati

Institut Jean-Nicod

Conference 'Les défis d'Hilary Putnam'

22-24 March 2005, Paris

Putnam is known for having demonstrated the existence of a new form of context-dependence, namely that which characterizes natural kind terms (and possibly others as well). Terms like 'tiger' and 'water' are indexical, Putnam says, since their conditions of application (i.e. the property something has to possess in order to be in the extension of the term) varies with the context of use — in a suitably broad sense of 'context'. A term like 'water' or 'tiger' is conventionally associated with a stereotype, i.e. a cluster of properties known to any competent user of the language. The stereotype enables one to identify paradigmatic exemplars of the category in one's local environment. But the property which determines whether or not something belongs to the extension of 'water' or 'tiger' is not the cluster in question : it's a different property *which the local exemplars happen to instantiate as a matter of empirical fact*. In our environment, the local exemplars that fit the tiger-stereotype are animals, but in a different environment, they would be robots controlled from Mars. In our context, therefore, 'tiger' refers to a certain type of animal, while in the counterfactual context the same word, associated with the same stereotype, would refer to a different type of entity sharing only superficial characteristics with *our* tigers.

The story is well-known, and I will not dwell on it. What I will focus on in this talk is the relation between Putnam's semantics and a body of views I call 'contextualism'. Contextualism *generalizes* context-sensitivity : it claims that sentences carry content only in the context of a speech act. Sentences, in vacuo, are not truth-evaluable. This view was put forward by ordinary language philosophers in the mid-twentieth century, and it has re-surfaced in recent times in the works of philosophers like John Searle, Charles Travis, and myself. I think Putnam's semantics has strong affinities with contextualism, and I want to argue this point.

A good place to start is Waismann's classic paper 'Verifiability', and more specifically the passage where he introduces the notion of 'open texture'.¹ That passage is representative of a view (and a type of example) which Waismann shared with Austin and Wittgenstein, and which has been revived by Travis and Searle. According to this view, truth-conditional content is essentially unstable and context-dependent. As we shall see, the thought-experiment conducted by Waismann in that passage is similar in some respects to Putnam's twin-Earth thought-experiment.

What Waismann says in introducing open texture can be paraphrased as follows. It seems that we have no problem assigning truth-conditions to ordinary sentences. To do so we describe a state of affairs, i.e. a type of situation, the obtaining of which is necessary and sufficient for the sentence to be true. That seems easy enough to do. Thus we think we know the truth-conditions of 'There is a cat next door', or 'This is gold' (or 'This is a man'): we can specify a state of affairs *s* such that the utterance is true if and only if *s* obtains. But this is an illusion — we can't really. Given an utterance *u* and the state of affairs *s* which is its alleged truth-maker, it is always possible to imagine a world in which *s* obtains, yet it is not the case that *u* is true (with respect to that world).

To show that, one has only to *embed* the state of affairs *s* within a larger situation, by providing further details about an imagined world in which *s* obtains. If the world in question is sufficiently unlike our world (for example, if the 'cat' talked about turns out to speak Latin, or grows to a fantastic size, or changes into a fish), we shall be at a loss when it comes to deciding whether the statement 'There is a cat next door' is actually true, with respect to that world, even though the state of affairs we initially specified obtains. Is a world in which there is an animal next door exactly like a cat in all respects except that it speaks Latin a world in which there is a *cat* next

¹ Waismann, Friedrich 'Verifiability', in Flew, Anthony (ed.), *Logic and Language*, 1st series (Blackwell, 1951), pp. 119-123.

door? We don't know, because "most of our empirical concepts are not delimited in all possible directions".² Or, to take another classic example:

The notion of gold seems to be defined with absolute precision, say by the spectrum of gold with its characteristic lines. Now what would you say if a substance was discovered that looked like gold, satisfied all the chemical tests for gold, whilst it emitted a new sort of radiation? 'But such things do not happen'. Quite so; but they *might* happen, and that is enough to show that we can never exclude altogether the possibility of some unforeseen situation arising in which we shall have to modify our definition. Try as we may, no concept is limited in such a way that there is no room for any doubt. We introduce a concept and limit it in some directions; for instance, we define gold in contrast to some other metals such as alloys. This suffices for our present needs, and we do not probe any farther. We tend to *overlook* the fact that there are always other directions in which the concept has not been defined. And if we did, we could easily imagine conditions which would necessitate new limitations. In short, it is not possible to define a concept like gold with absolute precision, i.e. in such a way that every nook and cranny is blocked against entry of doubt. That is what is meant by the open texture of a concept.³

We find something analogous to Waismann's embeddings in Searle's writings on the 'background'. In 'Literal Meaning', Searle enquires into the truth-conditions of 'The cat is on the mat'. It's easy to describe the sort of state of affairs that would make the sentence true (with respect to a particular assignment of values to indexical expressions). But once we have described such a state of affairs, we can embed it within an extraordinary situation:

Suppose that the cat and the mat are in exactly the relations depicted only they are both floating freely in outer space, perhaps outside the Milky Way galaxy altogether... Is the cat still on the mat? And was the earth's gravitational field one of the things depicted...?

² *Ibid.*, p. 120.

³ *Ibid.*

What I think is correct to say as a first approximation in answer to these questions is that the notion of the literal meaning of the sentence "The cat is on the mat" does not have a clear application, unless we make some further assumptions, in the case of cats and mats floating freely in outer space.⁴

Now Waismann has an explanation for the fact that, in extraordinary situations, 'words fail us', as Austin puts it. Here it goes:

If I had to describe the right hand of mine which I am now holding up, I may say different things of it: I may state its size, its shape, its colour, its tissue, the chemical compound of its bones, its cells, and perhaps add some more particulars; but however far I go, I shall never reach a point where my description will be completed: logically speaking, it is always possible to extend the description by adding some detail or other. Every description stretches, as it were, into a horizon of open possibilities: however far I go, I shall always carry this horizon with me. (...) [This] has a direct bearing on the open texture of concepts. A term is defined when the sort of situation is described in which it is to be used. Suppose for a moment that we were able to describe situations completely without omitting anything (as in chess), then we could produce an exhaustive list of all the circumstances in which the term is to be used so that nothing is left in doubt; in other words, we could construct a *complete definition*, i.e. a thought model which anticipates and settles once and for all every possible question of usage. As, in fact, we can never eliminate the possibility of some unforeseen factor emerging, we can never be quite sure that we have included in our definition everything that should be included, and thus the process of defining and refining an idea will go on without ever reaching a final stage.⁵

Waismann's point is simple enough. A term is used in, or applies to, situations of a certain type — situations with which we are acquainted. (Or at least: our mastery of the term is somehow connected to our having experienced the type of situation in

⁴ Searle, John 'Literal Meaning', in *Erkenntnis* 13 (1978), p. 211.

⁵ Waismann, 'Verifiability', pp. 121-3.

question.) To define a term is to describe the type of situation in question. The problem is that the situations in question, like Waismann's right hand or any aspect of empirical reality, cannot be completely described: they possess an indefinite number of features, some of which may never have been noticed and perhaps will never be noticed by anyone. When we describe an empirical situation, we make certain features explicit, but an indefinite number of other features remain implicit and constitute a sort of hidden 'background'. (Here I use Searle's term, on purpose.) Now the applicability of a term to novel situations depends on their similarity to the source situations, i.e. to the situations by association with which the term has acquired the meaning it has. For the term to be (clearly) applicable, the target situation must be similar to the source situations not only with respect to those features which easily come to mind and constitute the 'explicit' definition of the term (the 'tip', to use the iceberg metaphor), but also with respect to the hidden background. If the two situations considerably diverge with respect to the latter, it's unclear whether or not the term will be applicable, even though the explicit conditions of satisfaction are satisfied.

Other examples can be provided where the divergence between the source situation and the target situation affects not only the hidden background but also the explicit part (the tip of the iceberg). Thus Searle imagines that 'Snow is white' is uttered as a description of the following situation:

Suppose that by some fantastic change in the course of nature the earth is hit by an astronomical shower of radiation that affects all existing and future water molecules in such a way that in their crystalline form they reflect a different wave length when in sunlight from what they did prior to the radiation shower. Suppose also that the same shower affects the human visual apparatus and its genetic basis so that snow crystals look exactly as they did before. Physicists after the shower assure us that if we could see snow the way we did before, it would look chartreuse but because of the change in our retinas, which affects our observation of snow and nothing else, snow looks the same color as ever and will continue to do so to ensuing generations... Would we say that snow was still white?⁶

⁶ 'The Background of Meaning', p. 230.

Putnam's twin-Earth thought-experiment has the same structure.⁷ It's a case of divergence between source situations and target situation, where the divergence affects the tip and not merely the hidden background. For if we are to define water, we shall say that it's a liquid with such and such phenomenal properties, and that its chemical structure is H₂O, in the same way in which the whiteness of snow has to do both with its reflecting a certain wavelength *and* with its looking to us a certain way. We can imagine an extraordinary situation where the two things that go together as part of the tip are actually divorced from one another ; e.g. a situation in which the liquid which has all the superficial characteristics of water is not H₂O but XYZ. In such cases we don't know what to say because there is some measure of similarity between the source situations and the target situation, but we have no contextual clue as to which dimension of similarity matters.

That is not the conclusion Putnam himself draws from his thought-experiment. According to Putnam, twater is definitely *not* water. Instead of saying that *we don't know* whether or not twater counts as water, as the contextualists would do, Putnam claims that we *do* know : necessarily, since water is H₂O, twater is not water. I will come back to this difference between Putnam and the contextualists in a moment. First, however, I want to mention a couple of other respects under which Waismann's position diverges from Putnam's.

II

Waismann's argument for the instability of application conditions can be objected to on two grounds. First, it presupposes 'descriptivism', i.e. the idea that the only way to define words is to do so descriptively. Second, it conflates the semantic and the epistemological. In these two respects, Waismann's position seems to be dramatically different from that of contemporary philosophers of language like Putnam and Kripke.

Let me start with the first objection. All that Waismann's argument establishes, arguably, is that words can't be defined *in purely descriptive terms*. But this does not

⁷ Putnam, Hilary 'The Meaning of 'Meaning'', in his *Philosophical Papers*, vol. 2 : *Mind, Language and Reality* (Cambridge University Press, 1975), pp. 223-7.

imply that words cannot be defined *at all*, or that they do not possess definite (stable) conditions of application. The reason why words can't be defined in purely descriptive terms is that they have an irreducibly referential dimension: they, as it were, point to real situations in the world — situations which have an indefinite number of features and cannot be exhaustively described. Why not simply incorporate this referential dimension into our statements of truth- and application-conditions? Why not follow Putnam and define a cat as an animal belonging to the same species as *this specimen* (or those specimens, where the specimens in question are normal cats found in the local environment)? Why not define gold as *this metal* (while pointing to a piece of gold)? By explicitly incorporating ostensive reference to the actual environment into the definition of predicates like 'gold' and 'cat', as Putnam's semantics does, it seems that we can overcome the alleged unstability of truth-conditions. 'There is a cat next door' will be true, in the imagined situation (where a catlike animal next door speaks Latin), provided the catlike animal belongs to the same species as *our* cats (the cats to be found in ordinary situations).

(Second objection) Of course, we do not *know* whether or not the strange animal we imagine would be considered a cat by the scientists, just as we don't know whether or not a gold-like metal emitting a new sort of radiation would count as gold. But our epistemic limitations do not prevent the word 'gold' (or the word 'cat') from having a definite content and a definite extension: gold is anything that is the same metal as this, and a cat is any animal of the same species as those. Our epistemic limitations, so much emphasized by Waismann, have no bearing on the properly semantic issue.

These objections, inspired by Putnam's work, are well-taken but they are not decisive. In particular, they do not threaten Waismann's conclusion regarding the unstability of truth- and application-conditions. They would threaten that conclusion if, by appealing to ostensive definitions, we could determine a stable content (i.e. stable conditions of application) for the word thus defined. But we cannot. For a stable content to be determined by an ostensive definition à la Putnam, *the dimension of similarity to the demonstrated exemplars must itself be fixed*.⁸ But what fixes the

⁸ As McKay and Stern point out, "to determine the extension of the natural kind term, we need, in addition to the sample, at least some indication of the breadth of the term

dimension of similarity? On Waismann's picture, words are associated with situations of use, that is all. To apply the word to or in a novel situation, that situation must be similar to the source situations; but we cannot survey in advance all the possible dimensions of similarity between the source situations and possible target situations: open texture again.

According to Putnam (*op. cit.*), the predicate 'water' means something like: 'same-L as the transparent, odorless, thirst-quenching stuff to be found in lakes and rivers in the local environment', or more simply 'same-L as *that*' (pointing to a sample of water). If we change the demonstrated stuff (i.e. the liquid which satisfies the stereotype of water in the local environment) we thereby change the extension of 'water'. 'Water', therefore, has an indexical component, Putnam says. Yet there is another form of context-sensitivity in play here, which Putnam's ostensive definitions do not properly capture: *the dimension of similarity itself is not given, but contextually determined*. In some contexts the chemical composition of the demonstrated stuff will be relevant, in other contexts only functional properties will be relevant.⁹ Accordingly XYZ will count as water in some contexts simply because it's that liquid which is odorless, colorless, quenches thirst, and can be found in lakes and rivers. If the conversation bears on the issue of what currently fills a certain bottle, milk or water,

— the respects in which other individuals [in the extension] must be related to the sample' (McKay, Thomas and Stern, Cindy 'Natural Kind Terms and Standards of Membership', in *Linguistics and Philosophy* 3 [1979], p. 27).

⁹ A reader of Saul Kripke's *Naming and Necessity* (Blackwell, 1980) might object as follows :We are concerned only with what words mean in our language, not with what they mean in other possible situations ; now in our language water is H₂O, and 'This is water' is true iff the demonstrated stuff is H₂O. — This pseudo-Kripkean objection is mistaken and easy to rebut. I am talking about *our* language. Our language is such that we can imagine contexts in which it would be true to say 'This is water' of some stuff if the stuff in question had certain phenomenal/functional properties, whether or not the thing in question was H₂O.

then plainly the answer 'It's water' will be true, even if it's XYZ rather than H₂O.¹⁰

Putnam himself expresses awareness of this point:

x bears the relation same-L to y just in case (1) x and y are both liquids, and (2) x and y agree in important physical properties.... Importance is an interest-relative notion. Normally the 'important' properties of a liquid or solid, etc., are the ones that are *structurally* important; the ones that specify what the liquid or solid, etc., is ultimately made out of... From this point of view the important characteristic of a typical bit of water is consisting of H₂O. But it may or may not be important that there are impurities... And structure may sometimes be unimportant; thus one may sometimes refer to XYZ as water if one is *using* it as water...

Even senses that are so far out that they have to be regarded as a bit 'deviant' may bear a definite relation to the core sense. For example I might say 'did you see the lemon', meaning the *plastic* lemon. A less deviant case is this: we discover 'tigers' on Mars. That is, they look just like tigers, but they have a silicon-based chemistry instead of a carbon-based chemistry... Are Martian 'tigers' tigers? It depends on the context.¹¹

We now have in hand all the ingredients which I take to be constitutive of contextualism: the demonstrative component of meaning; the central role of similarity in determining extension (hence truth-conditions); and finally the context-dependence of similarity relations.

III

On the contextualist picture, words are not primitively associated with abstract 'conditions of application', constituting their conventional meaning (as on the Fregean picture). The conditions of application for words must be contextually determined, like

¹⁰ Chomsky, Noam *New Horizons in the Study of Language and Mind* (Cambridge University Press, 2000), p. 41.

¹¹ 'The Meaning of 'Meaning'', pp. 238-9.

the reference of indexicals. What words, qua linguistic types, are associated with are not abstract conditions of application, but rather *particular applications*.

In the spirit of Wittgenstein, consider what it is for someone to learn a predicate P. The learner, whom I'll call Tom, observes the application of P in a particular situation S; he associates P and S. At this stage, the semantic potential of P for Tom is the fact that P is applicable to S. In a new situation S', Tom will judge that P applies only if he finds that S' sufficiently resembles S. To be sure, it is possible for S' to resemble S in a way that is not pertinent for the application of P. The application of P to S' will then be judged faulty by the community, who will correct Tom. The learning phase for Tom consists in noting a sufficient number of situations which, like S, legitimate the application of P, as opposed to those, like S', which do not legitimate it. The semantic potential of P for Tom at the end of his learning phase can thus be thought of as *a collection of legitimate situations of application*; that is, a collection of situations such that the members of the community agree that P applies in or to those situations. The situations in question are the *source-situations*. The future applications of P will be underpinned, in Tom's usage, by the judgement that the situation of application (or *target-situation*) is similar to the source-situations.

In this theory the semantic potential of P is *a collection of source-situations*, and the conditions of application of P in a given use, involving a given target-situation S'', are *a set of features which S'' must possess to be similar to the source-situations*. The set of features in question, and so the conditions of application for P, will not be the same for all uses; it is going to depend, among other things, on the target-situation. One target-situation can be similar to the source-situations in certain respects and another target-situation can be similar to them in different respects. But the contextual variability of the conditions of application does not end there. Even once the target-situation is fixed, the relevant dimensions for evaluating the similarity between that situation and the source-situations remain underdetermined: those dimensions will vary as a function of the subject of conversation, the concerns of the speech participants, etc.

One particularly important factor in the contextual variation is the relevant 'contrast set'. As Tversky has pointed out, judgements of similarity are very much

affected by variations along that dimension.¹² If we ask which country, Sweden or Hungary, most resembles Austria (without specifying the relevant dimension of similarity), the answer will depend on the set of countries considered. If that set includes not just Sweden, Hungary and Austria but also Poland, then Sweden will be judged more like Austria than Hungary; but if the last of the four countries considered is Norway and not Poland, then it is Hungary which will be judged more like Austria than Sweden. The explanation for that fact is simple. Poland and Hungary have certain salient geopolitical features in common which can serve as basis for the classification: Hungary and Poland are then put together and opposed to Austria and Sweden. If we replace Poland by Norway in the contrast set a new principle of classification emerges, based on the salient features shared by Norway and Sweden: in this new classification Hungary and Austria are back together. Tversky concludes that judgements of similarity appeal to features having a high 'diagnostic value' (or classificatory significance), and that the diagnostic value of features itself depends on the available contrast set.

So the set of similarity features on which sense depends itself depends upon the relevant contrast set, and the relevant contrast set depends upon the current interests of the conversational participants. It follows that one can, by simply shifting the background interests ascribed to the conversational participants, change the truth-conditions of a given utterance, even though the facts (including the target-situation) don't change, and the semantic values of indexicals remain fixed. Charles Travis has produced dozens of examples of this phenomenon of truth-conditional shiftiness over the last thirty years, and his examples often involve manipulating the relevant contrast-set.¹³

¹² See Tversky, Amos 'Features of Similarity', in *Psychological Review* 84 (1977), 327-52.

¹³ See Travis, Charles *Saying and Understanding* (Blackwell, 1975), *The True and the False* (Benjamins, 1981), *The Uses of Sense* (Clarendon Press, 1989), and *Unshadowed Thought* (Harvard University Press, 2000). The following example, inspired from Austin, is taken almost at random from a list of Travis-examples compiled by Claudia Bianchi (then a graduate student of mine):

In this framework the background-dependence of truth-conditions emphasized by Searle is accounted for by appealing to the *global* character of the similarity between target-situation and source-situations. As Waismann stresses the source-situations are concrete situations with an indefinite number of features. Some of these features are ubiquitous and their diagnostic value in a normal situation is vanishing.¹⁴ They belong to the most general and immutable aspects of our experience of the world: gravity, the fact that food is ingested via the mouth, etc. When we specify the truth conditions of a sentence (for example the sentence 'The cat is on the mat'), or the conditions of application of a predicate (for example the predicate 'on' in that sentence), we only mention a small number of features — the 'foreground' features — because we take most of the others for granted; so we do not mention gravity, we presuppose it. Nevertheless, gravity is one of the features possessed by the situations which are at the source of the predicate 'on'; and there is an indefinite number of such features. These background features of the source-

Fred is walking with his young nephew beside a pond where a decoy duck is floating. Pointing to the decoy, he says, "That's a duck". Again we might ask whether what he said is true or false. But again, the above description is not enough for us to tell. If Fred has just finished laughing at a sportsman who blasted a decoy out of the pond, and if he has been trying to show his nephew how to avoid similar mistakes, then what he said is false. But suppose that Fred and his nephew are attending the annual national decoy exhibition, and the boy has been having trouble distinguishing ducks from geese. Then what Fred said may well be true. It would also be true had Fred said what he did in pointing out the fact that all the other ducks were poor copies (perhaps on the order of Donald Duck). (*Saying and Understanding*, p. 51)

¹⁴ See Tversky, 'Features of similarity', p. 342: "The feature 'real' has no diagnostic value in the set of actual animals since it is shared by all actual animals and hence cannot be used to classify them. This feature, however, acquires considerable diagnostic value if the object set is extended to include legendary animals, such as a centaur, a mermaid or a phoenix."

situations can be ignored inasmuch as they are shared by the situations of which we may wish to speak when we utter the sentence; but if we imagine a target-situation where the normal conditions of experience are suspended, and where certain background features of the source-situations are not present, then we shatter the global similarity between the target-situation and the source-situations. *Even if the target-situation has all the foreground features which seem to enter into the 'definition' of a predicate P, it suffices to suspend a certain number of background features in order to jeopardize the application of P to the target-situation.* That shows that the semantic potential of P is not, as in Fregean semantics, a set of conditions of application determined once and for all, but a collection of source-situations such that P applies to a target-situation if and only if it is relevantly similar to the source-situations.

A caveat: as Searle himself emphasizes, the fact that the target-situation does *not* possess certain background features of the source-situations does not automatically entail the non-applicability of the predicate P. It can be that the background features which the target-situation does not possess (for example gravity) are contextually irrelevant and do not affect the application conditions of the predicate.¹⁵ For the same sort of reason, the possession by the target-situation of what I have called the foreground features of the source-situations is no more a necessary condition for the application of the predicate than it is a sufficient condition. For a predicate (or a sentence) to apply to a target-situation that situation must resemble the source-situations under the contextually relevant aspects. So a predicate can apply even if the target-situation differs markedly from the source-situations, as long as, in the context and taking into account the contrast set, the

¹⁵ It is easy to imagine a context with respect to which it would be definitely true to say 'the cat is on the mat' of the gravitationless situation described by Searle:

For example, as we are strapped in the seats of our spaceship in outer space we see a series of cat-mat pairs floating past our window. Oddly, they come in only two attitudes. From our point of view they are either as depicted in Figure 1 or as would be depicted if figure 1 were upside down. "Which is it now?", I ask. "The cat is on the mat", you answer. Have you not said exactly and literally what you meant? (Searle, 'Literal Meaning', p. 212.)

similarities are more significant than the differences. Thus, in certain contexts, as Putnam notices, the predicate 'lemon' will apply to plastic lemons, or the word 'water' to XYZ.

I conclude that Putnam's position, suitably radicalized, leads directly to contextualism. The important contextualist idea — which forms the core of Austin's theory of truth — is twofold: first, words are associated with worldly situations (at whatever level of abstractness and fine-grainedness), i.e. entities at the level of reference rather than the level of sense; second, the sense of words, i.e. their conditions of application, depends upon similarities between those situations and the target-situation in which or to which the word is applied. Both ideas are present in Putnam's semantics. The only thing that has to be added to get contextualism is a third idea which is also (though less centrally) present in Putnam, viz. the idea that the relevant similarity relations themselves are not fixed once for all but depend upon the context.

Works cited

- Austin, John 'Truth', in his *Philosophical Papers* (2nd edition, Oxford University Press, 1970), 119-33.
- Chomsky, Noam *New Horizons in the Study of Language and Mind*, Cambridge University Press, 2000.
- Kripke, Saul *Naming and Necessity*, Blackwell, 1980.
- McKay, Thomas and Stern, Cindy 'Natural Kind Terms and Standards of Membership', in *Linguistics and Philosophy* 3 (1979), 27-34.
- Putnam, Hilary 'The Meaning of 'Meaning'', in his *Philosophical Papers, vol. 2 : Mind, Language and Reality* (Cambridge University Press, 1975), 215-71.
- Recanati, François *Literal Meaning*, Cambridge University Press, 2004.
- Searle, John 'Literal Meaning', in *Erkenntnis* 13 (1978), 207-24.
- Searle, John 'The Background of Meaning', in Searle, John, Kiefer, Ferenc, and Bierwisch, Manfred (eds.) *Speech Act Theory and Pragmatics* (Reidel, 1980), 221-32.
- Travis, Charles *Saying and Understanding*, Blackwell, 1975.
- Travis, Charles *The True and the False: : the Domain of the Pragmatic*, Benjamins, 1981.

Travis, Charles *The Uses of Sense: Wittgenstein's Philosophy of Language*, Clarendon Press, 1989.

Travis, Charles *Unshadowed Thought*, Harvard University Press, 2000.

Tversky, Amos 'Features of Similarity', in *Psychological Review* 84 (1977), 327-52.

Waismann, Friedrich 'Verifiability', in Flew, Anthony (ed.) *Logic and Language*, 1st series (Blackwell, 1951), 117-44.