

ISSUES IN PRAGMATICS (PLIN M301) 2006-07

LEXICAL PRAGMATICS

8. Attributive uses of concepts

1. Introduction

In the last two weeks, we've looked at two ways in which the concept expressed by use of a word on a given occasion may differ from the concept encoded: (a) **lexical narrowing**, where the concept expressed is more specific than the concept encoded; (b) **lexical broadening**, where the concept expressed is more general than the concept encoded. Today I want to look at a still further type of case: (c) **attributive use**, where a word or phrase is used to represent a concept whose content the speaker need not fully understand or endorse, but wants to attribute to someone else (or to herself at another time). [NB, for those of you who have looked at Donnellan's referential-attributive distinction, don't confuse today's notion of 'attributive use of a concept' with Donnellan's notion of 'attributive use of a NP' – I'll make a few comments on the relations between the two at the end of the lecture.]

I will try to show that attributive uses of concepts (and more generally **interpretive** or **metarepresentational** uses) play a major role in lexical acquisition, lexical borrowing and lexical-semantic change. I'll argue, too, that they can help with the analysis of what philosophers have called **deferential uses of concepts**, or **division of linguistic labour**, which have been seen as raising problems for the internalist approach to language and word meaning favoured by most Chomskyan linguists. Let me start with the philosophical problem, and then move on to more standard cases.

2. Division of linguistic labour

I've mentioned several times this term that even fluent speakers may have difficulty telling which objects fall under their concepts. Here are some examples:

- (1a) We may feel we understand the words *elm* and *beech*, and know they refer to different types of tree, but be unable to pick out an elm or a beech with any certainty;
- (1b) We may feel we understand the word *measles*, and know it's a type of disease, but be unable to diagnose a case of measles.
- (1c) We may know that *pewter* is a type of metal, but be unable to distinguish it from other, similar metals.
- (1d) We may know that *ketch* and *yawl* are types of ship, but be unable to tell them apart.

As we saw in Lecture 4, section 4, philosophers like Putnam (1975) have used facts like these as evidence that we may have only partial knowledge of the meanings of the words we use. According to Putnam, in many cases we may have a **stereotype** or a **prototype** – bits and pieces of encyclopaedic and perceptual knowledge – which are good enough for everyday classificatory purposes, but which need not provide a fool-proof method of deciding which objects actually fall under our concepts and which don't. (In one type of case, we have a **detector**, and we've seen that detectors aren't foolproof, cf. the fly and the little black dot from Lecture 4. In another type of case, we may not even have encountered the object, and don't have a detector at all.) We don't need to know more, says Putnam: we can leave the details to experts. This is what philosophers call **division of linguistic labour**, or **deferential** use of concepts: we defer to experts when it comes to deciding what our words ultimately refer to, or what they mean.

In fact, some philosophers have made the stronger claim that people may not just be **ignorant** but actually **mistaken** about what their words mean. Tyler Burge (1979) discusses a case involving the word *arthritis*. Imagine someone who mistakenly thinks that arthritis is not a rheumatic disease that occurs in the joints, but a much more general sort of disease. This person says things like (2):

(2) I've got arthritis in my thigh.

When asked what he thinks *arthritis* means, he says it means any general rheumatic pain. A Chomskyan might say that in this individual's idiolect, *arthritis* simply means 'pain', and the utterance is therefore true. According to Burge and other philosophers, the meaning of *arthritis* isn't under the individual's control: it is fixed by experts, and the utterance is simply false. Evidence that our use of terms such as *arthritis* is deferential is that when our mistake is pointed out, we don't normally say 'Well, that's what *arthritis* means in **my** idiolect'; we say, 'Oh, I was wrong: it wasn't arthritis at all.'

As I mentioned in Lecture 4, these arguments have been seen as presenting a challenge to Chomsky's internalist view of language. Recall that Chomsky (in *Knowledge of Language*, chap 2) draws a distinction between **E-language** (externalised language) and **I-language** (internalised language). An E-language is something like English, seen as a set of sentences existing independently of the mind of any given individual; individuals merely have partial knowledge of it, and may be mistaken in their beliefs about it. An I-language, by contrast, is

the language represented in a given individual's mind. Different individuals have different I-languages, which are **determined** by the individual's mental state: it makes no sense to talk of the individual having only partial knowledge of his own I-language. For Chomsky, E-language is not a well-defined object: for him, the most basic notion is I-language, and its study falls within the domain of individual psychology.

If we translate Chomsky's argument from the level of whole sentences to the level of individual words and concepts, it suggests a distinction between **I-concepts** (internalised concepts) and **E-concepts** (externalised concepts), and a conclusion that I-concepts are the proper object of linguistic study. Yet what Putnam, Burge and other philosophers are claiming is that the meanings of many words, e.g. those in (1a-d) above, are E-concepts, of whose contents the individual may have only partial knowledge. This raises a question: how can we reconcile Chomsky's internalist approach to language with cases such as (1a-d) above, in which it does seem that the individual may have only partial knowledge of word meanings, and defers to experts about what they mean? (Burge 1979, incidentally, endorses Chomsky's internalist approach to I-language in general, but rejects the internalist approach to concepts, for the reasons just given.)

In Lecture 4, I introduced a partial answer, showing how knowledge of word meanings does not require full understanding of the metaphysical truth conditions of the concepts they encode. In the next section, I'll introduce the notion of **attributive** (and more generally **interpretive**, or **metarepresentational**) use of concepts, and show how it may help to analyse a further range of linguistic phenomena, from concept acquisition to lexical-semantic change, where we may be seen as having only partial knowledge of word meanings. At the end of the lecture, I'll return to the 'deferential use of concepts', and show how the notion of attributive use might also help with this.

3. Attributive use of concepts

In *Relevance*, apart from lexical narrowing and broadening, we consider a third type of case which can lead to discrepancies between the concept encoded and the concept expressed. In **echoic** or **attributive** uses, the speaker uses language not to represent her own views, but to represent views she attributes to someone else. In **echoic** uses of language, the speaker not only represents views she attributes to someone else, but simultaneously expresses her own attitude to the attributed content (see *Relevance* chapter 4, sections 7-9). In this section, I'll survey a range of examples which might be analysed as involving attributive or echoic use not of a

whole utterance or proposition, but of some sub-part of it, e.g. a word or concept.

(a) Irony

The best known example of echoic use of concepts is in verbal irony. You're all familiar with the relevance-theoretic claim that irony is a case of echoic use, in which the speaker implicitly dissociates herself from an attributed utterance or thought. An example is (3), which may be seen as an echo of a particular utterance by a weather forecaster, or of a more general disappointed hope:

(3) It's a lovely day.

In much of the published relevance-theoretic literature, irony is presented as involving the echoic use (with expression of a dissociative attitude) of a whole proposition. It's easy to see, though, that a single concept may also be echoically used. Consider the following exchange:

(4a) *Jane*: Look at that sweet little doggie.

(4b) *Mary*: That 'sweet little doggie' is the terror of the neighbourhood.

Here Mary echoes and dissociates herself from the content of Jane's description 'sweet little doggie'. Notice that the utterance as a whole is descriptively used: Mary asserts that the dog in question is the terror of the neighbourhood, and commits herself to the truth of this assertion. Only the description *sweet little doggie* is used echoically, with a dissociative attitude, and this is responsible for the element of irony in (4b). We might think of the phrase *sweet little doggie*, and the concept it expresses, as occurring within quotation marks, so that (4b) expresses something like the proposition in (5):

(5) That 'sweet little doggie' (as you – absurdly – call it) is the terror of the neighbourhood.

Understanding Mary's utterance in (4b) would involve turning a regular (descriptively-used) concept into an **attributive concept** by the addition of something like quotation marks, and recognising that Mary was quoting and dissociating herself from a description attributed to someone else (here, Jane). Attributive concepts may be seen as a type of ad hoc concept. As in (4b), the current speaker takes no responsibility for the truth of the descriptive content of this ad hoc concept, but passes this responsibility on to someone else.

(b) Metalinguistic negation

A second example of the echoic use of concepts with dissociative attitude is the metalinguistic

negation in (6):

(6a). *Peter*: Are you happy to have finished your exams?

(6b) *Mary*: I'm not happy, I'm ecstatic.

If Mary were using the word *happy* in the regular way, her utterance should be a contradiction, because ECSTATIC entails HAPPY. In fact, she is echoing and dissociating herself from Peter's use of the word, not because it is false, but because it is inadequate to express her true feelings:

(7) I'm not merely 'happy' (as you – inadequately - call it), I'm ecstatic.

Again, the utterance as a whole is descriptively used, but contains a single expression that is echoically used, and from whose content the speaker explicitly dissociates herself. (See Carston 2002 chapter 4, section 4.4.2-4.4.4 for further discussion of metalinguistic use. See Wilson 2000 section 4 for the relation between metalinguistic uses and other types of echoic/attribution use.)

(c) Echoic allusion

Not all echoic uses of concepts involve a dissociative attitude. Consider (8):

(8) His life was a constant struggle against the slings and arrows of outrageous fortune.

Here there is an echoic allusion to be picked up. A hearer who notices it will understand the speaker as expressing something like the proposition in (9):

(9) His life was a constant struggle against 'the slings and arrows of outrageous fortune' (as Shakespeare called them).

Here, the speaker is not dissociating herself from but endorsing the attributed description, and since the attitude is one of endorsement, the hearer won't go too far astray if he fails to pick up the allusion; however, he will lose whatever cognitive effects the link to Shakespeare or Hamlet might yield (e.g. he will lose any implications that depend on the suggestion that the person referred to has properties in common with Hamlet). (For some reason, sports commentators are very keen on such examples: for instance 'Steffi Graf will not be there to praise the American but to bury her', alluding to the speech from *Julius Caesar*, 'I come not to bury Caesar but to praise him'.)

(d) Concept acquisition

Dan Sperber and I argued in the draft of *Relevance* (1983, dept file) that the attributive use of concepts might play a significant role in vocabulary acquisition. The child constantly hears utterances containing unfamiliar words, some of which will encode concepts he has not yet acquired (because he has not yet encountered any gerbils). As a result, he will be unable to represent the full proposition expressed by the utterance that contains them. Still, he may well understand enough to recover at least some of its logical and contextual implications. How can this be done?

Suppose, for example, a child hears his parents use the word 'gerbil' for the first time in an utterance like (10):

(10) Gerbils make good pets.

He would like to store this information in memory and use it derive whatever implications he can – for example, that gerbils are pets, that his parents may be thinking of buying a gerbil; to ask questions and form hypotheses about gerbils; to discover what a gerbil looks like, and so on. In other words, he would like to develop a fully adequate descriptive concept to associate with the word, and in the meantime do at least some partial processing of (10). A solution to both these problems might be to construct an attributive concept.

This might involve opening a conceptual address “GERBIL”, in quotation marks, with the lexical entry *gerbil*, a tentative logical entry containing the meaning postulate GERBIL →KIND OF ANIMAL, and an encyclopaedic entry containing the information that gerbils make good pets, that his parents may be thinking of buying one, and so on. None of this would enable the child to recognise gerbils, or to form an adequate idea of what a gerbil is, and in this sense his concept is defective. But the word *gerbil* could still be used to represent the non-defective concept he attributes to his parents. In this way, he will be able to refer correctly to the class of gerbils, to think and talk about them, to commit himself to whatever other people commit themselves to by their use of the word, and so on, while waiting for his own fully descriptive concept to develop. The ability to use concepts in this way, to represent some fuller concept attributed to others, is thus a valuable tool for the processing of defective conceptual representations.

This ability, of course, is not limited to children. We all come across new vocabulary items all the time. For example, suppose I hear you say that you're going to an auction hoping to buy a Louis Quinze escritoire. I have no idea what an escritoire is, though I assume it's a

piece of furniture; I couldn't recognise an *escritoire*, or know what to use it for, and there's no reason to think I have the concept in any full sense. However, when you come back, I ask you:

(11) Did you get the *escritoire*?

The best way to describe this is to say that I am using an attributive concept, intending to refer to whatever you were referring to when you used it. The proposition expressed would be something like (12):

(12) Did you get the '*escritoire*' (as you called it)?

(e) Concept change

Dan Sperber and I also argued in the draft of *Relevance* that the notion of attributive use of concepts might help with the analysis of lexical-semantic change. Take the French words *aperitif* and *digestif*, which seem to have undergone a parallel semantic change. Originally, *aperitif* described any drink that stimulated the appetite and *digestif* described any drink that aided digestion. The word *aperitif* now describes alcoholic drinks taken before a meal and *digestif* describes alcoholic drinks taken after a meal; the idea that they are stimulants to appetite or digestion has been lost. How did this change take place? We suggested that at an intermediate stage the original concept began to be attributively used: used, that is, to refer to drinks that are **claimed** to stimulate the appetite or aid digestion, but with no endorsement of the accompanying claims. With the words thus emptied of their original descriptive content, new referential and descriptive content could be gradually introduced, and the words would begin to be used in this new descriptive sense.

(f) Concept borrowing

We also argued that lexical borrowing from other cultures might be analysed along similar lines. Take the word *chic*, which in French has quite a wide range of meanings, e.g. 'skilful', 'sympathetic', 'elegant', etc., but in English has been narrowed down to meaning 'elegant, fashionable'. Originally, English users of this word would have been in the same position as the child acquiring an unfamiliar term: they would have a concept with a lexical entry and a few items of encyclopaedic information, but nothing amounting to a full descriptive concept. However, this concept could still be used attributively, to refer to whatever its original users use it to refer to, with or without an endorsement of the accompanying descriptive claims. Arguably, it has now been naturalised and acquired a new descriptive content of its own.

(g) Permanently attributive concepts

If this line of argument is correct, and attributive uses of concepts are often necessary precursors to full descriptive concepts, it seems reasonable to ask whether, for some of us at least, certain concepts never complete their development to full descriptive status. This might be either because we never fully discover their descriptive content, or because we refuse to endorse it when we do. For example, there must be many people who know the legal terms *tort*, *misdemeanour* and *felony*, and who know that they are used to refer to illegal acts, but not what type of acts they refer to. These people might be described as having attributive rather than descriptive concepts of a felony, a tort, etc. These cases bear obvious similarities to those mentioned by Putnam under the heading 'division of linguistic labour', where we defer to experts about the meanings of our words. There are certainly people who for much of their lives use such concepts as ANGEL, GHOST or CLAIRVOYANT echoically, while dissociating themselves from the associated descriptive claims: e.g., who use *clairvoyant* to refer not to those who can really foretell the future but to those who **claim** (or are claimed) to be able to foretell the future. Almost any concept, in any utterance, might have this extra quotational element, leading to discernible differences in implications.

It's quite conceivable, in fact, that there are certain concepts which we all use attributively or echoically all the time, and which no-one fully understands at all. Certain slang terms, like *tacky*, or *gross*, presumably fall into this category: we all know roughly what they mean, but there's no expert who knows exactly what they mean, and no metaphysical truth conditions to be discovered. More abstract cultural and religious terms like *love* or *heritage* might work in a similar way: each of us uses them to refer to whatever other people, those other people we feel in sympathy with, refer to when they use them; these people in turn use them to refer to whatever other people they are in sympathy with use them to refer to, and so on indefinitely. When we come to inquire what all these uses have in common, it may well turn out that they are cases of what Wittgenstein called **family resemblance terms** (see Lecture 3), and their meanings are genuinely fuzzy. However, their existence would be linked to our ability to use concepts attributively, and would do nothing to undermine the claim that many concepts do have well-defined descriptive meanings with either cognitively or metaphysically necessary and/or sufficient conditions attached. We should thus be able to do justice both to intuitions of fuzziness and to the fact that many terms are quite precisely defined.

4. Philosophical applications of the notion of attributively used concept

Let's return now to the Burge/Putnam cases involving **division of linguistic labour** or

deferential uses of concepts. We might investigate the possibility of reanalysing at least some of them as cases of attributive use, as with the examples of *tort*, *misdemeanour* and *felony* mentioned above. This fits well with Burge's comment (1979: 79) that:

'The argument [about social determinants of lexical meaning, e.g. of *arthritis*] can get under way in any case where it is intuitively possible to attribute a mental state or event whose content involves a notion that the subject incompletely understands...'

My suggestion is that at least some cases of 'incomplete understanding' should be analysed in terms of attributive use. This should make it possible to reconcile Chomsky's claim that we should concentrate on **I-concepts** with Putnam's or Burge's claim that in some cases we may not know what our own words or concepts refer to, and leave their meaning to someone else. Genuinely deferential uses of concepts work in just this way. This approach might also handle cases where there are no experts, and the meanings of our concepts are genuinely fuzzy or undergoing change (Margolis 1998 suggests that it can also be used in the analysis of concept-acquisition).

However, there are other cases which do not seem to be genuinely deferential in the sense of Burge and Putnam, but which can be adequately handled in one of the ways I've suggested today. Consider (13), a quote from George W. Bush when he was running for President:

(13) There was quite a riff-raff between us over our future health policy.

Here, by *riff-raff* Bush presumably meant something like 'argy-bargy', and (13) might well be true on that interpretation. By contrast, if we treat this as a genuinely deferential use, since *riff-raff* actually means 'yobbish crowd', the utterance barely makes sense and would come out as false. Most of us would probably agree that here the Chomskyan internalist account based on speaker's idiosyncratic meaning works best: it is not that the speaker incompletely understands what he means, and would defer to others and agree that he meant 'There was quite a yobbish crowd between us...' – he knows what he means, and commits himself only to that. A similar account would work for many of Burge's other examples. (On the internalist account, *riff-raff* as Bush uses it might be either a wrongly-analysed descriptive concept or a partially-analysed interpretive concept whose analysis is incompatible with most people mean by that word.)

A second difference between deferential and attributive uses of concepts is that deferential uses of language seem to be necessarily accompanied by an **endorsing** attitude,

whereas, as we've seen, echoic uses of concepts may be accompanied by **questioning**, **dissociative** or merely **neutral** attitudes. This makes it possible to analyse some of Donnellan's supposedly **referential** uses of language as echoic, even though they are not deferential (the speaker rejects, rather than accepts, the attributed meaning). Suppose Peter and Mary see someone at a party drinking from a martini glass, and the following exchange occurs:

(14a) *Peter*: Who's the man with the martini?

(14b) *Mary*: The man with the martini is Jane's husband.

In fact, Mary knows that the man Peter intends to refer to is drinking water, but she still uses the phrase *the man with the martini* as the easiest way to pick him out for Peter. According to Donnellan, in this case *the man with the martini* is being **referentially** used. In referential use, the descriptive content of a concept does not contribute to the proposition expressed by an utterance), and the speaker is not necessarily committed to its truth. The analysis I'd suggest is to treat (14b) as a case of echoic use, in which Mary echoes Peter's expression without endorsing its descriptive content, as the optimally relevant way of picking the man out. By contrast, (14b) is not easily analysed as a deferential use of language, precisely because Mary does not **endorse** Peter's use of the phrase.

Finally, there are cases where an expression is used to represent or interpret concept which does not yet exist, and which is therefore not attributed to any user to whom the current speaker can defer. This often happens in developing a theory: for example, the word *relevance* in early versions of relevance theory was used to represent the full-fledged theoretical concept that we hoped to develop (and still do...). This is a case of **interpretive** use (based on resemblance of content), but not of **deferential** or **attributive** use, because there was no expert to whom we could defer, or whose views we were interpreting.

The notion of attributive – and more generally interpretive or metarepresentational – use of concepts thus seems to have a fair range of possible applications in both linguistics and philosophy. (Dan Sperber has explored some of these ideas further in his 1997 paper 'Intuitive and reflective beliefs; see also Richard Horsey 2000 on 'Meaning postulates and deference'.)

Homework

1. Make sure you understand the differences between **attributive**, **echoic** and **interpretive** uses of concepts.

2. Try to think of some of your own examples of attributive, echoic or interpretive uses of concepts falling into the various categories I've listed.

Reading

Wilson, D. 2000 Metarepresentation in linguistic communication, section 4. In D. Sperber (ed.) *Metarepresentations: A Multidisciplinary Perspective*. OUP, Oxford: 411-48.
(downloadable from my website: www.phon.ucl.ac.uk/home/deirdre/home.html)

Background references

- Putnam, H. 1975 The meaning of "meaning". Reprinted in Putnam, *Philosophical Papers*, vol 2.
- Burge, T. 1979 Individualism and the mental. In P. French, T. Uehling & H. Wettstein (eds) *Midwest Studies in Philosophy*, vol. IV.
- Carston, R. 1996 Metalinguistic negation and echoic use. *Journal of Pragmatics* 25: 309-330.
- Carston, R. 2002 *Thoughts and Utterances*. Chap 4, sects 4.4.2-4.4.4.
- Sperber, D. 1997 Intuitive and reflective beliefs. *Mind & Language* 12: 67-83.
- Sperber, D. & D. Wilson 1983 Draft of *Relevance* chap VIII, sect 4; sect 5, pp 64-71 (dept file)
- Recanati, F. 2000 *Oratio obliqua, oratio recta: The Semantics of Metarepresentations*. Part VI, Deference and metarepresentation. MIT Press. (I have a copy if you'd like one.)
- Margolis, E. 1998 How to acquire a concept. *Mind & Language* 13. Also in Margolis & Laurence 1998. (some discussion of deferential concepts).
- Horsey, R. 2000 Meaning postulates and deference. *UCL Working Papers in Linguistics* 12.
- Sperber, D. & D. Wilson 1998 Irony & relevance: A reply to Seto, Hamamoto & Yamanashi. In R. Carston & S. Uchida (eds) 1998 *Relevance Theory: Applications and Implications*. John Benjamins: Amsterdam. (& from my website)