

15. The Pragmatics of Deferred Interpretation

GEOFFREY NUNBERG

Subject	Theoretical Linguistics » Pragmatics
DOI:	10.1111/b.9780631225485.2005.00017.x

1 Deferred Interpretation

By deferred interpretation (or “deference”) I mean the phenomenon whereby expressions can be used to refer to something that isn't explicitly included in the conventional denotation of that expression. The interest in these phenomena stretches back to Aristotelian discussions of metaphor, and while the study of the mechanisms of deference has made considerable progress in the interval, that (literally) classical framework still underlies a lot of the assumptions that people bring to the phenomena. So it will be useful to address some of these legacies from the outset.

Traditional approaches tend to regard figuration (and by extension, deference in general) as an essentially marked or playful use of language, which is associated with a pronounced stylistic effect. For linguistic purposes, however, there is no reason for assigning a special place to deferred uses that are stylistically notable – the sorts of usages that people sometimes qualify with a phrase like “figuratively speaking.” There is no important linguistic difference between using *redcoat* to refer to a British soldier and using *suit* to refer to a corporate executive (as in, “A couple of suits stopped by to talk about the new products”). What creates the stylistic effect of the latter is not the mechanism that generates it, but the marked background assumptions that license it – here, the playful presupposition that certain executives are better classified by their attire than by their function. Those differences have an undoubted cultural interest, but they don't have any bearing on the more pedestrian question of how such usages arise in the first place.¹

This assumption about the stylistic role of figuration is closely linked to a second assumption of traditional approaches, the idea that deference is exclusively a pragmatic phenomenon. For example, Grice (1967) treats metaphor as a kind of conversational implicature that arises from a violation of the maxim of quality; on his view, metaphorical utterances invariably have a literal reading to which a truth value (usually, false) can be assigned and which constitutes the input to some inferential schema that generates a “secondary” figurative reading. The assumption is that deference is somehow inconsistent with conventionalization, so that we can say that a word has distinct lexical meanings only when the connections that once licensed its multiple uses have been somehow obscured or forgotten. That is what often leads people to characterize polysemy in diachronic terms and to talk about figurative meanings that have been lexicalized as “dead” or “frozen” metaphors. As Ravin and Leacock (2000) put it: “polysemes are etymologically and therefore semantically related, and typically originate from metaphorical usage.” But statements like this are better thought of as origin myths than as analytic hypotheses.² In fact, our assumption that such-and-such a usage is lexicalized is very often based on no more than an intuitive sense of its stylistic effects or an observation of its frequency, rather than on any strict analytical criteria.

But from a linguistic point of view, there's no reason to distinguish between the mechanisms that operate within the lexicon to produce meaning extensions and those that operate in a purely pragmatic way. Figuration doesn't necessarily cease to be figurative just because it is subject to some

conventionalized restrictions. What leads us to say that the processes that produce multiple uses of expressions are lexicalized is not that they are no longer transparent (they may very well be), but only that the language constrains or enriches their use over and above what could be predicted on pragmatic grounds alone. Conventionalization should not be confused with absolute arbitrariness; it makes more sense to think of deference as a process that is orthogonal (or more accurately, heterogonal) to the pragmatic mechanisms that give rise to deferred readings, analogous to other productive derivational processes. And conversely, the mere fact that a particular usage is both frequent and stylistically unremarkable doesn't necessarily mean that it is lexicalized (even if that criterion may lead lexicographers to include it in their dictionaries).

A further problem with the traditional view of deference is in the way it classifies the deferred uses of expressions, according to the conceptual relations or correspondences that they manifest. Synecdoche, for example, is defined by the third edition of the *American Heritage Dictionary*, in part, as "a figure of speech in which a part is used for the whole." Over and above the obvious category mistake here – what the dictionary means, of course, is that synecdoche involves using the *names* of parts in place of the *names* of wholes – there are reasons for keeping these relations distinct from the purely linguistic mechanisms that exploit them. For one thing, a single mechanism may exploit several distinct figures. There may be no purely linguistic reason, for example, for distinguishing a traditional synecdoche like *blade* for "sword" from a metonymy like *crown* for "monarch" or a metaphor like *wolf* for "rapacious person."³ And conversely, a single conceptual correspondence might figure in two distinct kinds of deferred interpretation. For example, the perceived relation between newspaper publishers and their products makes possible two interpretations of the objects in (1) and (2):

(1) Murdoch bought a newspaper last week.

(1) (pointing at a newspaper) Murdoch bought that last week.

Still, there are reasons for believing that (1) and (2) involve different linguistic mechanisms, the first affecting the use of descriptive terms and the second the use of demonstratives and indexicals. Finally, we will see that linguistic mechanisms of transfer are subject to certain constraints which aren't necessarily implicit in the conceptual relations they depend on, but which require the introduction of independent principles.

2 Meaning Transfers

With this as background, we can turn to the linguistic mechanisms that license the deferred uses of expressions. In this article, I will concentrate on the mechanism I will refer to as MEANING TRANSFER, which underlies what we ordinarily describe as the metaphorical and metonymic uses of names and descriptions. I will start by discussing meaning transfer as a purely pragmatic process, then turn to the way it is implicated in various lexicalized rules and schemas, and then finally discuss its application to some long-standing questions in syntax.

Meaning transfer is the process that allows us to use an expression that denotes one property as the name of another property, provided there is a salient functional relation between the two.⁴ These relations can obtain in virtue of a direct correspondence between properties, when one property calls up another that it resembles (as in metaphor) or evokes (as in synaesthesias like *a blue mood*). When we use the word *horseshoe* to refer to a logical operator shaped like a horseshoe, for example, we exploit a relation that can be characterized without reference to the circumstances of any particular horseshoes or any particular typographical marks. Or the relation can be mediated by relations between the bearers of the properties. This is what underlies transfers involving metonymy and synecdoche, such as when we use the word *novel* or the name of a particular novel to refer to the film rights to a work, as in "Spielberg bought the novel for \$1 million." In that case we exploit a correspondence that holds between distinct instances of film rights and distinct novels: there is exactly one of the former for each of the latter. In what follows I will be mostly talking about the second sort of transfer, but everything I say will apply with appropriate modifications to transfers of the first kind as well.

Meaning transfers can apply to predicates of any kind, whether lexical or phrasal, and whether used

attributively or predicatively. By way of developing some of the features of the process, let's consider (3):

(3) I am parked out back.

One might be tempted to say that the transfer in (3) applies to the subject *I*, in a sort of “driver for car” metonymy. But there are a number of reasons for assuming that the transfer here applies to the conventional meaning of the predicate. For example, if the speaker has two cars, he wouldn't say:

(4) We are parked out back.

though of course this would be an appropriate utterance if there were two people who were waiting for the car.⁵ Note, moreover, that we can conjoin any other predicate that describes the speaker, but not always one that literally describes the car:

(5) I am parked out back and have been waiting for 15 minutes.

(6) *I am parked out back and may not start.

For both these reasons, we assume that the predicate *parked out back* in (3) carries a transferred sense, which contributes a property of persons whose cars are parked out back.

Meaning transfer operates not just on the meanings of predicates or verb phrases, but on the meanings of common nouns, as well, whether they appear in predicate position or referring position. Take (7), as uttered by a restaurant waiter:

(7) Who is the ham sandwich?

The process of transfer is straightforward here; from the point of view of the waiter, at least, customers acquire their most usefully distinctive properties in virtue of their relations to the dishes they order. But in this case, unlike the “parked out back” examples, the relevant property is expressed by a common noun, which can equally well be used as the content of an NP in referential position in a sentence like (8):

(8) The ham sandwich is at table seven.

In (8), the predicate *ham sandwich* has a transferred meaning, where it contributes a property of people who have ordered ham sandwiches.⁶

3 Conditions on Meaning Transfer

As I noted earlier, meaning transfer is possible when there is a salient correspondence between the properties of one thing and the properties of another, in which case the name of the first property can be used to refer to the second. With an utterance like *I am parked out back*, for example, we begin with a functional correspondence between the locations of cars in a lot and the properties of the owners or drivers of these cars. When two property domains correspond in an interesting or useful way – of which more in a moment – we can schematize the operation of predicate transfer as follows:

(9) Condition on Meaning Transfer

Let P and P' be sets of properties that are related by a salient function $g_t: P \rightarrow P'$. Then if F is a predicate that denotes a property $P \in P$, there is also a predicate F' , spelled like F , that denotes the property P' , where $P' = g_t(P)$.⁷

A correspondence of this sort can hold in either of two cases. Sometimes there is a direct functional relation between two sets of properties, as in cases of metaphor and synaesthesia – for example in the relation between grades of temperature (cold, cool, warm, hot) and the affects they bring to mind.

In other cases, though, the correspondences between properties are mediated by correspondences between their bearers, which is what underlies metonymic and synecdochic transfers. There is no direct correspondence between the property of being parked out back and the distinguishing property of any particular person, save via the relation between a person and the thing that has that property. We can represent this particular case of meaning transfer as follows:

(10) Metonymic Transfers

Let h be a salient function from a set of things A to another (disjoint) set of things B . Then for any predicate F that denotes a property P that applies to something in A , we can represent the meaning of a derived predicate F' , spelled like F , as in either a or b:

- a. $\lambda P. \lambda y (\forall x_{[\text{dom } h]}. h(x) = y \rightarrow P(x))$
- b. $\lambda P. \lambda y (\exists x_{[\text{dom } h]}. h(x) = y \text{ and } P(x))$

Note that this entails that predicates of this type are in fact ambiguous between “universal” and “existential” readings, depending on whether all or only some of the bearers of the original property are in the inverse image of h for a given value. And in fact both types of reading are generally available. In cases like “I am parked out back,” we would normally assume that the speaker means to say that all the (relevant) cars he is looking for are parked out back, as in (10a).⁸ By contrast, when a painter says, “I am in the Whitney,” she doesn’t imply that all her paintings or even all her relevant paintings are in the Whitney, but only that something she painted is in the Whitney, as in (10b). And when an accountant says of her firm, “We are in Chicago,” she might intend either interpretation, depending on whether she’s talking about all of the firm’s offices or merely about one of them. Still, it is more useful to think of these two types of readings as two ways of instantiating the general schema given in (9), rather than as two distinct conditions that license predicate transfer. Meaning transfer is a single linguistic process.

4 The Criterion of Noteworthiness

The schemas in (9) and (10) do a reasonable job of representing the truth conditions associated with utterances like “I am parked out back” and “I’m in the Whitney,” but they miss some important pragmatic conditions on the use of such utterances. For example, suppose my car was once driven by Yogi Berra. Then according to the conditions in (10), I should be able to use the name of this property to describe the property that I acquire in virtue of my relationship to my car. But it would be odd for me to say:

- (11) ?I was once driven by Yogi Berra.

even in a context in which it might be relevant to say, “My car was once driven by Yogi Berra.” By the same token, a painter might say with reference to one of her paintings, “I’m in the Whitney Museum,” but not, ordinarily:

- (12) ?I’m in the second crate on the right.

Intuitively, the difference is this: when a painting goes into a museum its creator acquires a significant or notable property, whereas when it goes into a crate she doesn’t, at least not usually.⁹

Let me describe this condition by saying that predicate transfer is only possible when the property contributed by the new predicate is “noteworthy,” which is to say one that is useful for classifying or identifying its bearer relative to the conversational interests. In this sense noteworthiness is equivalent to what Downing (1977) meant when she said that novel noun–noun compounds must be “appropriately classificatory,” and to the conditions that Clark and Clark (1979) observed on the zero–derivation of English verbs from nouns. The fact that the criterion is applicable here demonstrates that the transfer process creates new predicates with new meanings, just as other derivational processes do.¹⁰

It is important to bear in mind that noteworthiness is not the same thing as relevance, though it is clearly a related notion. In this connection, consider (13) and (14), adapted from Jackendoff (1992):

(13) Ringo was hit in the fender by a truck when he was momentarily distracted by a motorcycle.

(14) ?Ringo was hit in the fender by a truck two days after he died.

Let's assume that these utterances exemplify transfers of the meanings of the relevant relational expressions – that is, that *Ringo* denotes the singer rather than his car.¹¹ The difference between the two cases is that when a truck hits Ringo's car while he is driving it, the event will probably have important consequences for him as well: he is likely to have been startled, or annoyed, or put to trouble and expense. Whereas once Ringo is dead, the things that happen to his car don't generally invest him with any properties worth mentioning.¹² But while the distinction is intuitively clear, our ability to characterize it formally requires that we be able to distinguish between the relevance of a proposition (e.g. that Ringo's car was hit) and the relevance of its trivial entailments (e.g. that Ringo has the property of having had his car hit). It may be that a suitable version of relevance theory will be able to clarify this distinction, but for the present purposes we can just take noteworthiness in an intuitive way.¹³

5 Predicate Transfer in Systematic Polysemy

The availability of transfer for common nouns, adjectives, and other lexical categories is what underlies the patterns of lexical alternation that have been described using such terms as “regular polysemy” (Apresjan 1973), “deferred reference” (Nunberg 1979), “semantic transfer rules” (Leech 1974), “sense transfer” (Sag 1981), “connectors” (Fauconnier 1985), “sense extensions” and “logical metonymies” (Pustejovsky 1991, 1995, Copestake and Briscoe 1995), “lexical networks” (Norvig and Lakoff 1987), “subregularities” (Wilensky 1991), and “lexical implication rules” (Ostler and Atkins 1991) – not to mention just plain “metonymy” and “metaphor” (see, among many others, Lakoff 1987). Needless to say, there are many differences among these approaches, both formally and in the data they are invoked to explain. But they all involve the same type of generalizations, which can be phrased as implicational statements of the form: “If an expression has a use of type U, it also has a use of type U'.” For example, the name of a writer can be used to refer to his or her works, a word that denotes a periodical publication or kind of periodical publication can be used to refer to the organization that publishes it; and a word that denotes a kind of plant or animal can refer to its meat or substance (this latter is the rule called “grinding”).

(15) Proust is on the top shelf.

(16) The *Chronicle* (the newspaper) opposed the highway project.

(17) We were eating chicken on tables made of oak.

Many of these rules are much more general in their application than the examples we have been discussing, and require no specialized context to license them. The correspondence between the properties of dishes and customers provides a useful means of identification only in the domain of a restaurant, and then only relative to the interests of waiters – we could think of usages like these as examples of Clark and Clark's (1979) “contextual expressions.” In general, we don't think of these as involving lexical senses of the items in question: their number is too open-ended, and their use tends to be restricted to particular types of contexts or subcommunities of speakers.¹⁴

But the property correspondences that license the transfers in (15)–(17) hold across a wider range of situations, and provide a more context-independent way of classifying the bearers of derived properties. In these cases we may very well want to say that the transferred predicate represents a lexical sense of the item in question, or at least deserves listing in a dictionary.¹⁵ To a certain extent, this is a relative matter. For example standard dictionaries often assign the word *white* a sense like “In chess, the person playing white.” They do this because even though the correspondence between a color and a role is context-specific, the derived predicates *white* and *black* are much more generally useful for classifying chess players than the property of having ordered a ham sandwich as a means of

classifying restaurant customers, since so many relevant things follow from which color a player takes.

As usages become progressively more useful and less specialized, we come to the more general patterns of lexical alternation that are commonly described by means of the formula "x for y," as in "artist for work" or the alternation sometimes described as "portioning," as in "I spilled a couple of beers." In the extreme case we can talk about very general patterns of alternation like the systematic polysemy of abstract nouns like *obstinacy* or *vanity*, which can refer either to a quality or to the extent of its instantiation in a particular case (see Aronoff 1976):

(18) Obstinacy is usually a mistake.

(19) Her obstinacy surprised us.

6 Semantics or Pragmatics?

At this point we can ask in what ways the alternations created by meaning transfer should be treated as semantic (that is to say, conventionalized), rather than being derived by purely pragmatic processes. As I noted earlier, this is not a simple matter of a privative opposition: the mere fact that an alternation is in some way conventionalized does not mean that it no longer has any pragmatic basis. In the extreme case, of course, alternations may be preserved in the lexicon as the *disjecta membra* of transfers that no longer have any productive role in the language, whether because the background assumptions that originally licensed them are no longer valid, or because the items have acquired specialized meanings over the years. We may still perceive a certain relation among the uses of *cell* to refer to small rooms, organizational units, and the structural units of biology, but no one would seriously propose that these could all be generated from some core sense via purely pragmatic inferences – over the years the senses have been enriched with lexically specific material (a political cell, for example, is necessarily clandestine). At the other end of the scale, there are uses like those in the *ham sandwich* cases that are obviously extralexical – that is, as generated exclusively by pragmatic principles, with no need of any lexical specification. In the middle, however, lies a very broad range of productive alternations whose status is less clear.

There are several factors that militate for a semantic treatment of certain alternations. In the first place, alternations are sometimes subject to language-specific restrictions that don't seem to have any obvious pragmatic explanations. For some reason or another, for example, the English rule of grinding seems not to apply to yield the names of liquids:

(20) ?We fried the chicken in safflower (olive, corn, etc.).

(21) ?I enjoy a glass of orange (pear, apple, etc.) with my breakfast.

The meaning that *safflower* would have in (20) is clear enough – what else would you fry things in, if not an oil? – and it may very well be that in some kitchens cooks routinely use *safflower* in this way, but this is not the general practice. Analogously, in English we use the names of painters to refer to their works, and of designers to refer to their clothes, but we don't do this with the names of composers, and we do it with the names of novelists and filmmakers only when they produce genre fiction or genre films:

(22) She owns two Picassos and a Renoir.

(23) The sale racks at the store were full of Jill Sanders.

(24) ?We listened to a lovely Scarlatti (?Steven Sondheim).

(25) I love to curl up with a good Agatha Christie (Simenon, ?Italo Calvino, ?Dostoyevsky).

(26) One of my favorite Hitchcocks (?Renoirs) is playing at the Bijou.

There may very well be some subtle ethnoaesthetic explanation for these patterns, but it is by no means obvious what it could be – certainly it can't be argued that the interpretation of *a Dostoyevsky* would be any harder to figure out than the interpretation of *an Agatha Christie*. On the face of things,

this is just a fact of English.¹⁶

The rule of portioning provides similar examples. As Copestake and Briscoe (1995) observe, the rule is highly productive: we can say, "I drank two beers," "I drank two Michelobs," and so forth. But for some reason the rule doesn't apply to the names of wines. "I drank two Sauternes last night" can only involve reference to two types of wine, not two glasses. It may be that this difference has a connection with a difference in the way beer and wine were historically sold, but in the wine-bar-saturated context of modern London or San Francisco, it certainly feels like an arbitrary constraint.

We can make a similar point about the use of words like *cheerful* and *sad* to apply to places or circumstances that evoke certain emotional responses. We speak of *a cheerful room* or *a sad turn of affairs*, but the alternation doesn't extend to items like *afraid* and *frightened*; we don't speak of *an afraid house*, for example. Pustejovsky (1995) has noted that this behavior correlates with the fact that sad-type predicates do not take prepositional objects except by adjunction (e.g. *sad about the loss*), whereas passive participles like *frightened* do (e.g. *frightened of the bear*). The observation seems sound, but it is hard to see how we could get from this syntactic property to a purely pragmatic explanation of which predicates allow this alternation.

Similar patterns emerge when we look at systematic polysemy cross-linguistically. For example, there seems to be no salient cultural reason why French-speakers should use the names of fruits to refer to the brandies made from them (*une prune*, *une poire*) whereas English-speakers do not.¹⁷ Apresjan (1973) notes that in Russian you can use the name of an organ to denote a disease of that organ, as in "She has a kidneys." But this usage is not common in English, though it would be comprehensible enough.

More generally, languages can differ not just in the particular alternations that they permit, but in the general tolerance of polysemous processes. According to Jerrold Sadock (personal communication), Greenlandic Eskimo permits few types of systematic polysemy; for example, it does not allow grinding of animal names to produce mass terms for either meat or fur, though you can apply grinding to tree names to get terms for wood. So you can't say, "He eats walrus" or "He wears walrus," but you can say "His boat is made of oak." Here too, we seem to be dealing with arbitrary lexical restrictions, though in this case we would want to cast them as very general principles.

But even when there are strong arguments that militate for treating a given pattern of systematic polysemy as lexicalized, there may still be reasons for looking to pragmatics to explain the intricacies of its application. There are various exceptions to general patterns of polysemous alternation that can be explained only by reference to the noteworthiness criterion that I described earlier. For example, the name of a publication can be used to refer to a publisher only when the publication is one that is usually produced by an organization dedicated to that task – a newspaper or travel guide, say, but not a cookbook or Latin grammar. Only in those circumstances would there be an identificationally useful correspondence between the distinctive properties of the publication and those of its publisher. Clearly, though, this isn't the sort of property that would be specified in the lexical entries for items like *newspaper* and *cookbook*: the essential properties of those types of publications are independent of the particular economic structures in which they are produced.¹⁸ In the same way, while we saw that the application of the rule of grinding is subject to various lexical constraints, it answers to pragmatic considerations as well. We don't ordinarily use the names of breeds to refer to the meat of animals, for example, but we sometimes use them to refer to their hides or fleeces: we say *She wears angora* but not *We don't eat angora*. But this isn't a matter of lexical meanings, and it would take only a slight modification of culinary or sartorial practices to change the pattern.

A similar point could be made about the well-studied phenomena involving verb-class alternations. For example, Dowty (2000) has pointed out that the criterion of noteworthiness is relevant to determining when predicates permit the "swarm alternation" exemplified by (27) and (28):

(27) Bees are swarming in the garden.

(28) The garden is swarming with bees.

Dowty points out that like other derivational processes, such as noun-noun compounding, new predicates formed on the model of (28) are felicitous only if they contribute "the property a place or

space has when it is 'characterized' by an activity taking place within it – that is, when the extent, intensity, frequency, and/or perceptual salience of [the] activity [that] takes place there is sufficient to characterize the Location in a way that is relevant for some purpose in the current discourse" (2000: 122). So while there is no question that the *swarm* alternation is a lexicalized construction of English, its application is nonetheless subject to broadly pragmatic conditions on meaning transfer.

7 Transfer in Composition

The observations about noteworthiness that we've made here help to shed some light on compositional processes, as well. It's well known, for example, that adjectives can exhibit a considerable flexibility as regards the relation of their conventional meaning to their heads. Following Leech (1974), Fillmore (1978) has noted that the adjective *topless* can be used in any of the following ways:

- (29) a. *topless swimsuit*, "swimsuit without a top."
- b. *topless dancer*, "dancer who wears a swimsuit without a top."
- c. *topless bar*, "bar featuring dancers who wear swimsuits without a top."
- d. *topless district*, "district containing bars featuring dancers who wear swimsuits without a top."
- e. *topless legislation*, "legislation regulating bars featuring dancers who wear swimsuits without a top."

As Fillmore points out, there are two ways of analyzing examples like these: we can assume either that the adjective has a transferred meaning in (29b)–(29e) or that the adjective retains its literal meaning and the looser connection between it and the noun meaning is made possible constructionally in one way or another.

Pustejovsky (1991) opts for a constructional approach in connection with the analogous example of *a fast motorway*, which has a meaning equivalent to "a motorway permitting fast driving." On his analysis, this interpretation arises when the adjective is applied in its conventional meaning to an event embedded in the telic structure of the noun *motorway*, which yields an interpretation as in (30):

- (30) ... *Telic* = $\lambda e[\text{drive on (e, people) and fast (e)}]$

This line of analysis could in theory be applied to examples like:

- (31) *fast food*, "food intended for rapid service and consumption."
- (32) *free-range chicken (eggs)*, "chicken meat (eggs) derived from free-range chickens."

There are reasons, though, for thinking that the multiple readings here arise from lexical transfers rather than from constructional vagueness or ambiguity. Note that only certain adjectives can appear in these constructions:

- (33) A fast (?drunken) motorway.
- (34) Fast (?rude) food.
- (35) Free-range (?beheaded) chicken.
- (36) Topless (?Speedo) bars.

These observations can be explained by appeal to the criterion of note-worthiness: the property of motorways that permits fast driving is significant to classifying them, and for this reason we can derive a new adjective *fast'* to denote this property. But the property of motorways that permits drunken driving has no particular classificatory usefulness, and hence the conventional meaning of *drunken* does not correspond to a noteworthy property of motorways. Analogously, we can usefully classify foods according to whether they are characteristically eaten or served quickly, but not according to whether they are eaten or served in a rude manner. And while the way a chicken was

nourished corresponds to a noteworthy property of its meat, the fact that it was killed by beheading does not. Finally, we may find it useful to single out a class of bars whose dancers wear topless garments, but people to date have not demonstrated an interest in distinguishing classes of bars according to the brand of swimsuit that their dancers wear.¹⁹

Note that this is not to say that postulating a deferred meaning for an expression is always preferable to offering a constructional solution or to assigning the expression a more general monosemous meaning. For example, consider examples like (37)–(39):

(37) We enjoyed the movie.

(38) We enjoyed the book.

(39) We enjoyed the talk.

In principle, we could understand the objects in (37)–(39) as having transferred senses that enable them to refer, respectively, to the processes of watching a movie, reading a book, and listening to a talk. But there seems to be no criterion of noteworthiness associated with this construction – the object of *enjoy* can be just about any noun that can be associated with some kind of process that could evoke pleasure, from a radio to a thermos bottle to a printing press.²⁰ It seems preferable, then, to assume that the interpretations of (37)–(39) are supplied constructionally, possibly following the proposals of Pustejovsky (1991) and Copestake and Briscoe (1995).²¹

8 Noteworthiness as a Diagnostic

The examples we've been discussing show how the criterion of noteworthiness provides a useful diagnostic for determining whether and where meaning transfer is present. Consider sentence (40), from Copestake and Briscoe (1995):

(40) The south side of Cambridge voted Conservative.

On the face of things, we might analyze (40) in either of two ways: either the description within the subject NP has a transferred meaning that describes a group of people, or the VP has a transferred meaning in which it conveys the property that jurisdictions acquire in virtue of the voting behavior of their residents. In this connection, Briscoe and Copestake observe that not all geographical descriptions can be used in constructions like these:

(41) Three villages/three villages south of the river/?three villages built of stone voted for the proposed ban on timber production.

Briscoe and Copestake note that “it appears that only modifiers which might apply to the group of people, or which are locational (*the south side of Cambridge*) are fully acceptable” (1996: 45). We might put this by saying that geographical expressions can have transferred meanings in which they apply to groups of people only when the property they denote correlates usefully with a distinguishing characteristic of those groups. This property need not necessarily be locational – one might say, for example, *Villages with large numbers of detached single-family houses tend to vote Conservative*, since the kinds of housing a village has may correlate with some distinctive properties of people who live there.

But the restrictions on this type of transfer make clear that the transfers in (40) and (41) is on the interpretation of the subject NP. An analogous point could be made about a sentence like (42):

(42) The huge (?domed) stadium rose as one to cheer the team.

The size of a stadium correlates with an important property, the magnitude, of the group of people who fill it; its architectural features do not. Again, this supports the analysis of (42), in which the subject has a deferred interpretation.²²

Note, by contrast, that other words that denote places or physical locations can be used without restriction to refer to the people who frequent or inhabit them. Consider the difference between (43)–(44) and (45)–(46), after some examples suggested by Cruse (2000):

- (43) The factory is out on strike.
- (44) ?The factory rose as one to cheer the contract.
- (45) The school is taking a day off.
- (46) The school rose as one to cheer the football team's victory.

Factory behaves like *village* – it can have a transferred sense in which it applies to people just in case it provides a noteworthy property of the group. But *school* seems to have a single lexical meaning that allows it to denote both a building and the people who use or run it. Following Pustejovsky (1995), we could say that *school* has a “dot object” structure, which provides for its use to refer to things of different types, though nothing turns on this particular form of analysis.²³ But whether we can postulate a univocal analysis of a word depends on how and whether the criterion of noteworthiness comes into play in determining its uses.

9 Syntactic Consequences: “Sortal Crossings”

On the basis of the examples I gave by way of introducing the notion of predicate transfer, like *I am parked out back*, it might be thought that the process is generally confined to highly context-particular situations, particularly as it affects phrasal constituents. In fact, though, the mechanism is applied very widely and usually passes without notice. This point has a particular relevance to the syntactic phenomena that we can think of as “sortal crossings,” as in (47), suggested by Jackendoff, and (48):

- (47) Ringo squeezed himself into a narrow space.
- (48) Yeats did not like to hear himself read in an English accent.

On the standard way of thinking about these phenomena, the reflexives in these examples present a problem, since they seem not to be coreferential with their antecedents. In (47), the subject *Ringo* refers to a person, whereas the reflexive appears to refer to his car. In the heyday of imperial syntax, there were proposals to handle sentences like this one by means of syntacticosemantic operations like “car deletion” or with other formal devices of roughly equivalent effects, and other people have suggested various semantic and pragmatic approaches.

But approaches like these all have the same failing, in presuming that the reflexive is not strictly coreferential with the subject. In fact we might better think of a sentence like (47) as involving a transfer in which *squeeze into a narrow space* has a meaning that applies to persons – it denotes the relation that people enter into in virtue of the maneuvers they perform with their car. In this case the reflexive and its antecedent are coreferential in a strict sense. And by the same token, for (48) we can assume that people who read Yeats's poetry aloud are doing something to him as well. But here again, the transfer is only available when the derived property or relation provides some useful or noteworthy information about its bearer, in accord with modern assumptions about authorship. So we wouldn't ordinarily say:

- (49) ?Yeats wrote a lot of himself in sprung rhythm.

When Yeats writes a poem in a particular metrical pattern, that is, he is not performing an operation on himself, in the way that someone who reads his poetry aloud is also doing something to him.

All of this leads to a strong hypothesis: natural language permits no sortal crossings in any of the rules or constructions that ordinarily impose conditions of identity – not just with reflexives, but with pronouns, relative constructions, and so forth. In a sentence like (50), for example, we can assume that the content of the clause *was featured in a Madonna video* has a deferred meaning in which it

contributes the property that newspaper companies acquire in virtue of the exposure given their publications:

(50) The newspaper that Mary works for was featured in a Madonna video.

But not everything that happens to a newspaper copy confers a noteworthy property on its publisher, of course:

(51) ?The newspaper that Mary works for fell off the table.

Stallard (1993) suggests an analogous example:

(52) No airlines that fly to Denver are based on the East Coast.

In this case the predicate in the relative clause contributes the noteworthy property that an airline acquires from the activities of the flights it operates. But other activities of flights don't have these consequences for the airlines that operate them:

(53) ?The airline disappeared behind a mountain.

A related hypothesis would stipulate that there need be no relaxing of the coordinate structure constraint to deal with sentences like:

(54) Roth is Jewish and widely read.

Instead, we will interpret both *Jewish* and *widely read* as predicates that contribute properties of persons. In this way we can honor the intuitions that originally motivated semanticists to appeal to Zeugma as a way of determining whether a word has one or more senses.²⁴

ACKNOWLEDGMENTS

Versions of this material can be found in Nunberg (1979, 1995) and Nunberg and Zaenen (1992). For comments on these articles, I am grateful to Sue Atkins, Ted Briscoe, Cleo Condoravdi, Ann Copestake, Lauri Karttunen, James Pustejovsky, Ivan Sag, Tom Wasow, and Annie Zaenen. I am grateful to Gregory Ward for comments on this article.

1 The study of the cultural and cognitive implications of the belief systems that underlie various forms of deferred use has been an important theme in the program of cognitive linguistics; see, among many others, Johnson (1987), Lakoff (1987), Lakoff and Turner (1989), Sweetser (1990).

2 The view of polysemes and idioms as “dead metaphors” is mostly wishful, when you think about it, since there is no evidence that most such usages were ever “live” in the sense that they were transparent and optimal ways expressing certain meanings. Despite the earnest efforts of armchair philologists to find the “true story” behind figures like *pull the wool over someone's eyes* or *rob Peter to pay Paul*, there's no reason to believe that these expressions were ever completely transparent – that there was a time when everybody had a particular Peter and Paul in mind.

3 Another reason for distinguishing between types of figuration and the linguistic processes that exploit them is that the same relations that underlie metaphor, metonymy, and the rest can be used in non-linguistic systems of communication, such as the icons on a computer desktop.

4 I am speaking here of properties, but everything I say will generalize to relations and other types.

5 By the same token, in Italian the predicate adjective agrees with the subject, so in (i), with “my father” as subject, we would express the transferred meaning using a masculine adjective *parcheggiato* for “parked,” even though the word for “car” is a feminine, *la macchina*:

(i) Mio padre è parcheggiato (**parcheggiata*, fem. sing.) in dietro. “My father is parked out back.”

6 In some discussions, examples like (8) have been analyzed as involving a kind of “reference transfer” or “deferred reference,” with the implication that an actual ham sandwich must figure in the interpretation of the utterance – that is, that the transfer operates on the NP interpretation. But there are a number of reasons for concluding that the transfer here takes place on the common noun meaning – that is, that this is a case of meaning transfer, rather than reference transfer. One way to make this point is to consider the interpretation of the determiner in the phrase *the ham sandwich*, which doesn't presuppose the existence of a unique ham sandwich (think of a waiter in a fast-food restaurant who is standing in front of a table piled with ham sandwiches), but does presuppose the existence of a unique ham-sandwich orderer. Or we can consider some examples involving anaphor. Fauconnier (1985) gives examples (i) and (ii):

(i) The mushroom omelet was eating with chopsticks.

(ii) The mushroom omelet was eating itself/himself with chopsticks. The use of a reflexive in (ii) would presume that the object of *eat* was introduced by the subject NP –that is, that the subject NP actually referred to a mushroom omelet on the route to its ultimate interpretation. But inasmuch as the transfer actually takes place at the level of the common noun, which contributes only a property of persons, the example is ill formed. Gregory Ward (personal communication) has observed that it is in fact possible to say things like:

(iii) The ham sandwich is complaining because the bread is soggy. where it might be argued that it is only in virtue of the introduction of a ham sandwich as a discourse referent that we can infer the uniqueness of the bread. I agree that utterances like (iii) are more felicitous than the referentially equivalent (iv):

(iv) The customer at table 7 is complaining because the bread is soggy. It can be argued, though, that the relative felicity of (iv) is due more to the mode of presentation of the reference (i.e., as a ham-sandwich orderer) than in virtue of the introduction of any actual ham sandwich as a discourse referent. By analogy, consider how we sometimes take advantage of the nominal root of a derived adjective to establish the definiteness of some other thing that's associated with the reference of that root:

(v) I don't much like Italian food but I'd like to learn the language.

(vi) ?I don't much like the kind of food they serve at Mario's Grotto but I'd like to learn the language. In neither (v) nor (vi) is Italy itself an element of the discourse, but the country is made salient by the mode of presentation of a certain kind of food in (v).

7 The “salience” of a function depends on a number of factors. Among other things, the properties in the domain of the function (here, car locations) have to be discriminable, and the relation itself has to be manifestly familiar to participants. These conditions are schematized at some length in Nunberg (1995), though in that paper these transfer functions were defined over domains of individuals rather than properties, and no distinction was made between predicate transfer and what I call DEFERRED INDEXICAL REFERENCE, the operation that explains how you can point at a set of keys and say “That is parked out back.”

8 Ordinarily, of course, a single parking-lot patron is interested in only a single car, but we can imagine cases in which a single person has come with more than one car – or, more plausibly, we can take an author's utterance of a sentence like “I'm published by Knopf,” which on its most likely reading is equivalent to “All of my books are published by Knopf.”

9 Of course the noteworthiness of a property depends among other things on the conversational interests, and you could fiddle with the utterance or the context in such a way as to make most of these examples acceptable. A painter who feels she is being slighted in favor of other painters in her gallery, for example, might say: (i) Those daubers get one-person shows while I'm relegated to a crate in the basement. The most we can say, then, is that certain derived properties are canonically or stereotypically more noteworthy than others, a difference that will be important when we come to talk about the lexicalization of predicate transfer.

10 The requirement of noteworthiness is one feature that distinguishes meaning transfer from the process of deferred indexical reference that I mentioned in connection with (2), an utterance of *Murdoch bought that*, accompanied by demonstration of a newspaper. This follows from the fact that deferred indexical reference exploits correspondences between individual things, not the properties associated with lexical meanings – it does not create new predicates (or any predicates at all). For example, the name of a publication can't be used to refer to a publisher unless there is a salient functional relation between the distinctive properties of the first and the distinctive properties of the second, which usually obtains because the publishing organization is constituted to produce only that work. That explains why we can say *She works for a newspaper* but not *She works for a cookbook*. But if the identity of a publisher is evident in the

appearance of a particular cookbook, we can point at it and say *She works for them*. I will not take up the phenomenon of deferred indexical reference in this article; for a discussion, see Nunberg (1993).

11 Jackendoff analyzes these as involving transfer of the sense of *Ringo*, rather than of the relational expressions, but we've seen that there are a number of reasons for rejecting this approach.

12 Of course one could say, "Picasso had to wait until after he died to get into the Louvre." But in this case the transferred predicate here applies to the personage, who survives the person.

13 See Sperber and Wilson (1995) and the large body of literature on relevance theory that has grown out of their work.

14 From the linguistic point of view, perhaps, the most interesting forms of transfer involve the multiple uses of functors and grammatical categories – prepositions, for example, or simple verbs of motion. I won't have anything to say about these types of transfer here, simply because the subject is too vast.

15 This doesn't necessarily mean that the senses have been conventionalized. Dictionaries often assign separate senses to the common metaphorical uses of words, for example, even though these might be predictable. To take an example suggested by Fillmore and Atkins (2000), the fact that dictionaries assign the word *crawl* a sense "to act or behave in a servile manner" doesn't mean that people couldn't come up with this use of the word in the absence of a convention. Dictionaries are more concerned with recording what is conventional in the loose sense of the word (i.e. as opposed to what is unconventional) than in its strict sense (as opposed to what is non-conventional).

16 I would be happier with this conclusion if the same regularities didn't show up in more-or-less the same form in a number of other languages.

17 In French, one can use the expression *une menthe* to refer to a syrup made from mint, but one cannot use *une* (or *un*) *pomme* to refer to a juice made from apples. This last example shows that these are not simple cases of ellipsis of the head noun, which in the case of syrups (*un syrop*) would be masculine.

18 It could easily happen that cookbook publishing should become the province of dedicated organizations, as dictionary publishing is, in which case it would make sense to say "John works for a cookbook." But the lexical meaning of *cookbook* would not have changed in the process.

19 Though, as one of the editors points out to me, people do sometimes have an interest in distinguishing bars according to the material worn by their patrons – e.g. leather.

20 Copestake and Briscoe (1995) suggest that the sentence *We enjoyed the dictionary* is odd because dictionary is not lexically associated with a process. But it's easy to imagine a context in which this sentence could be used, and more generally, there is no requirement that the relevant process be explicit in the lexical meaning of the noun. When someone says, for example, *I enjoy the beach*, the understanding is that the speaker enjoys the activity of sitting around or walking at the beach, but there is nothing in the meaning of the word that specifies these processes.

21 There is no reason to assume that meaning transfer is involved in all processes of type-shifting, for example when a verb like *consider* coerces a predicate interpretation of its complement. In contrast to the cases we have been talking about, criteria of noteworthiness seem to play no role here: there is no constraint on what NP can be substituted for x in the frame "I consider it an x."

22 One thing we should note is that in some cases transfers can apply either to a subject or to a predicate in such a way that it can be difficult to tell which element has the transferred reading. In a sentence like *Stevens is challenging*, for example, we might assume either that a transfer occurs on the noun *Stevens* to enable it to contribute the property of a body of literary work, or that *challenging* has a transferred meaning that contributes a property that writers acquire in virtue of the properties of their works. Note that we could elaborate this sentence with either an animate or an inanimate pronoun:

(i) Many people find Stevens challenging, but we sell a lot of it.

(ii) Many people find Stevens challenging, but he regarded his poems as simple. In (i), transfer occurs on the noun *Stevens*; in (ii), on the predicate *challenging*.

23 Cruse would presumably explain this distinction by saying that *school* and *factory* are both monosemous, but would highlight different "facets" of lexical meaning. This is similar, I believe, to what

Langacker (1984, 1987) is getting at when he talks about "active zones," this in connection with examples like (i) and (ii):

(i) We heard the bugle.

(ii) I finally blinked. In (i), what we heard was **the sound** of a bugle; on Langacker's view this reading is created when the verb emphasizes a certain substructure of lexical meaning. In one form or another, this kind of analysis is certainly justified for many words that denote complex objects, but it is an empirical question whether we want to apply it in any given case, as opposed to postulating a meaning transfer. One reason for supposing that it may not be the right tack in (i) is that *bugle* can often designate a disembodied sound, as in:

(iii) The bugle floated faintly in the still night air. But this form of analysis certainly seems plausible for (ii), as an alternative to supposing a transfer from persons to eyelids. In the end, of course, the choice between a monosemous analysis and one involving meaning transfer can't always be cleanly resolved – there is no magic way to resolve the longstanding lexicographical debates between "splitters" and "lumpers" of word-senses. For an appreciation of just how complicated these issues can become when one tries to resolve them in full lexicographical detail, see Fillmore and Atkins (2000).

24 Note that sortal crossings of this sort are not possible when the transfers are metaphorical rather than metonymic; that is, when they involve a direct relationship between properties rather than one mediated by relations among their bearers. For example, we can't say (i):

(i) ?The second line of the proof begins with the horseshoe that's hanging on the wall of Deb's room. In this and similar cases, the particular things in the range of the transfer function don't acquire properties from anything that happens to the things in its domain – that is, nothing that happens to any real horseshoe has any consequences (noteworthy or otherwise) for any particular horseshoe-shaped symbol.

Cite this article

NUNBERG, GEOFFREY. "The Pragmatics of Deferred Interpretation." *The Handbook of Pragmatics*. Horn, Laurence R. and Gregory Ward (eds). Blackwell Publishing, 2005. Blackwell Reference Online. 28 December 2007 <http://www.blackwellreference.com/subscriber/tocnode?id=g9780631225485_chunk_g978063122548517>

Bibliographic Details

The Handbook of Pragmatics

Edited by: Laurence R. Horn And Gregory Ward

eISBN: 9780631225485

Print publication date: 2005