

SCALARS

Bart Geurts

Introduction

Scalar expressions are so-called because they line up in sequences like the following:

- (1) a. $\langle \textit{perfect}, \textit{good}, \textit{okay} \rangle$
b. $\langle \textit{boiling}, \textit{hot}, \textit{warm} \rangle$
c. $\langle \textit{all}, \textit{most}, \textit{many}, \textit{some} \rangle$
d. $\langle \textit{adore}, \textit{love}, \textit{like} \rangle$
e. $\langle \dots, \textit{three}, \textit{two}, \textit{one}, \textit{zero} \rangle$

Intuitively, *perfect* is better than *good*, which in its turn is better than *okay*, and therefore these are scalar expressions, which together constitute the scale in (1a). As the collection in (1) illustrates, the notion of scalar expression is not restricted by syntactic or semantic type.

There is a more or less standard theory of scalar expressions. It is the neo-Gricean account developed by Gazdar, Horn, Levinson, and others. I take it that this account is sufficiently well known that only a brief reminder is called for here.

- (2) a. Barney took some of the money.
b. Barney took all of the money.

According to the neo-Gricean theory, scalar *some* does not entail the negation of *all*, and therefore (2a) does not entail that (2b) is false (just as $\exists x\phi$ does not entail that $\forall x\phi$ is false). Nonetheless, (2a) would normally be taken to imply that (2b), but this is because in most cases (2a) licenses the conversational implicature that Barney didn't take all of the money. So a typical construal of (2a) has two components: its meaning or truth conditional content, which does not exclude that Barney took all of the money, and a conversational (or scalar) implicature, which does.

The neo-Gricean approach comes in several flavours. One distinction, which will loom large in this paper, concerns Grice's notion of generalized conversational implicature, or GCI. It is not quite clear what a GCI is supposed to be, but Grice's idea seems to have been that if a conversational implicature ϕ occurs often enough in the presence of an expression α , then the implicature will somehow become conventionally associated with α itself, although we (or rather Grice and his followers) would be unwilling to say that it has thereby become part of the meaning of α , or has turned into a conventional implicature.

It will be evident that scalar implicatures make for good candidate GCIs, and indeed many authors have argued that scalar implicatures are GCIs. This is the position taken by Horn, Gazdar, and Levinson, among others, although the latter is considerably more explicit on this score than his

comrades-in-arms. On the other hand, there are many authors who either prefer to ignore the distinction between generalized and particularized implicatures, or explicitly argue against it (e.g. Hirschberg 1985, Carston 1995). In the following, I will generally use the epithet ‘neo-Gricean’ to refer to those followers of Grice for whom the generalized/particularized distinction is of central importance, and who hold, in particular, that scalar implicatures are GCIs.

Prominent representatives of the neo-Gricean school, thus defined, are Horn, Levinson, and Gazdar, and the most explicit defense of it is Levinson (1997).

In this paper I will outline a theory of scalars which is only partly compatible with the neo-Gricean approach. I do believe that something like conversational implicature must be called upon to account for at least some uses of scalars, but I also believe that these uses are only part of a puzzle which is much more complicated than the neo-Griceans would have it. More importantly, however, I am sceptical about the notion of GCI, and am more congenial to theories that prefer to treat conversational implicatures as context dependent inferences. In a word, I am what Levinson calls a nonce-implicature theorist.

Before I proceed to expound and defend my position, let me define some terms that will expedite the following discussion. Take a possible scalar expression, the cardinal *two*, say. (Although I will argue later that cardinals are quite untypical scalars, if they deserve to be called ‘scalars’ at all, they are very convenient for expository purposes, because it is so easy to be precise about their possible meanings.) Supposing that *two* denotes a property of sets, I will say that the unilateral meaning of *two*, if it has one, is $\lambda x(\text{card}(x) \geq 2)$. The bilateral meaning of *two*, if it has one, is $\lambda x(\text{card}(x) = 2)$. To say that an expression has a unilateral or bilateral meaning is to say something about its truth-conditional import. This is not necessarily the same as saying something about its conventional, or lexical, meaning. As I understand my own terminology, it is not inconsistent to suggest, as I will do, that the lexical entry of a scalar only yields a unilateral meaning, although *in some contexts* it may actually have a bilateral meaning. If I say that a scalar licenses a unilateral or bilateral *interpretation*, then there is always a larger expression involved, usually a sentence, like:

(2) Fred ate two bananas.

On a given occasion, this may be interpreted as conveying that the number of bananas Fred ate was exactly two. Thus, the expression *two* licenses a bilateral interpretation of this sentence. The whole point of distinguishing between meaning and interpretation is that it allows us to say, if the spirit moves us, that a scalar expression has a unilateral meaning but may license a bilateral interpretation; this is how Horn (1989), like most neo-Griceans, analyzes (2). Of course, it is also possible that a scalar with a bilateral meaning licenses a bilateral interpretation; this is how Horn (1992) analyzes (2). It is even possible to say without inconsistency that a scalar with a bilateral meaning licenses a unilateral interpretation, on a given occasion. This may strike one as a purely theoretical possibility, but it is not (cf. Hirschberg 1985, Koenig 1991). The main claims that I will defend in this paper are the following:

- (i) Although all scalars can have bilateral meanings (in certain contexts),
- (ii) in general, the lexical meaning of a scalar expression is unilateral,
- (iii) but cardinals are an exception: their lexical meanings are bilateral.

None of these claims is original with me; it is rather the combination that is new. I will try to motivate it in the remainder of this paper.

Special contexts, special meanings

If I am not mistaken, there are contexts in which any scalar expression can have a bilateral interpretation. In such a context a scalar expression like *most*, for example, can be paraphrased as ‘most but not all’. In the following I will outline a defense of this claim which is a condensed version of an argument presented in more detail in Geurts (in press). But first I want to make it clear that this argument is not intended to pave the way for an account of how scalar expressions are interpreted *in general*. If my argument goes through, it merely shows that scalars have bilateral meanings under some, rather special, circumstances. It is possible that this is how scalar expressions are interpreted in general, but as we will see later, it is not likely that this is the case.

My argument is based on the behaviour of scalars under negation, as in the following examples:¹

- (3) a. Fred didn’t win many votes.
b. Fred didn’t win MANY votes: he won them ALL.
- (4) a. I don’t like this car.
b. I don’t LIKE this car: I LOVE it.

Normally speaking, an utterance of (3a) will be interpreted as entailing that Fred won less than many votes, which suggests rather strongly that the scalar expression *many* is being used with a unilateral meaning. However, (3b) is an exception: if in this case, too, *many* meant something like ‘many and possibly all’, then it would be a mystery how one can consistently deny that Fred didn’t win many votes and affirm that he won them all. The same holds, *mutatis mutandis*, for the pair of examples in (4). I call negated sentences like (3b) and (4b) ‘denials’.² I maintain that in the context of a denial a scalar expression can have a bilateral meaning. Thus in the context of (3b) *many* entails *not all* and in the context of (4b) *like* entails *not love*. This is not to say, however, that scalars are systematically ambiguous between unilateral and bilateral meanings. My proposal is, rather, that this is an instance of polysemy, and that in general the lexical meaning of a scalar expression is a unilateral one (cardinals being an exception, see below), which in special cases is restricted to a bilateral meaning. Most words are polysemous, which is to say that they have several senses that are related to each other. Familiar examples are:

¹ Working within the framework of relevance theory, Kempson (1986) and Carston (1996) have proposed analyses of scalar denials that are very similar to the one I endorse.

² My notion of denial is coextensive with van der Sandt’s (1991); it is approximately coextensive with Horn’s (1989) notion of ‘metalinguistic negation’. There are other varieties of denial beside the scalar cases discussed here, but these will not play a role in this paper.

- (5) a. Fred drank a lot of beer.
 b. Fred drank one beer.
- (6) a. This is my last Sony.
 b. Sony is a Japanese firm.

While in (5a) *beer* denotes a liquid, in (5b) the same word is used as a count noun; and while in (6a) *Sony* is used to refer to an electronic appliance, in (6b) it is the name of firm. Such patterns are numerous and ubiquitous, and at a pre-theoretical level it is fairly clear how they develop. It seems evident that the ‘liquid’ sense of *beer* is primary, but given that beer standardly comes in certain units, it is natural to refer these units as ‘beers’. Similarly, it is only natural that certain goods get their names from the firms that produce them, in this case Sony.

In the case of *beer* we may safely assume that both the mass and the count interpretation have become conventionalized, and the same holds, presumably, for the two senses of *Sony*.

However, conventionalization is not an essential ingredient of polysemy.³ Polysemy is a productive phenomenon, which can be context dependent to a very large degree. The following examples from Nunberg (1978) illustrate this:

- (7) a. He can hit the ball two *football fields*.
 b. He hit a home run two *games* ago.
 c. I love some kinds of liver; *chicken* is tasty.
 d. *The Times* asked if he could have an exclusive interview.

The italicized expressions in these examples have non-standard meanings, but it is evident that, for instance, (7c) can be understood even in the absence of a linguistic convention licensing the ‘chicken liver’ sense of *chicken*. On the other hand, there can be no doubt that what we observe here is basically the same phenomenon as in (5) and (6); the only difference is that in the earlier examples the relevant senses are probably conventionalized.

I assume that in all these cases the words in question mean what we intuitively take them to mean. Thus in (5b) the meaning of *beer* is ‘glass of beer’ (or ‘bottle of beer’, depending on the context), and in (7c) *chicken* means ‘chicken liver’. Although I do believe that this is very plausible, pre-theoretically speaking, it is not a position that it is universally accepted. In particular it might be held, as some philosophers have done, that on their literal meaning the examples in (7) are incoherent or worse, although they somehow manage to convey useful information. On this view, the noun *chicken* in (7c) can only express its ordinary sense, and therefore the sentence does not have a proper meaning, strictly speaking, although it may serve to communicate an intelligible message. I don’t find this position a particularly attractive one, and I believe it can be refuted, but I will not attempt to do so in this paper.

In all instances of polysemy, conventionalized or not, world knowledge is exploited to generate a derived sense from a another, more basic, sense. For example, if someone is acquainted with

³ Whereas in Geurts (in press) I use the term ‘semantic transfer’ to refer to non-conventionalized instances of polysemy only, and restrict my use of ‘polysemy’ accordingly, in this paper ‘polysemy’ does not entail conventionalization. I actually dislike both options, which may be due to the fact I have thus far failed to come to terms with concepts like conventionalization, conventional meaning, and so on.

the ‘liquid’ sense of *beer* and has the relevant world knowledge concerning beer, then there are only one or two plausible candidates for the count sense of beer. One way of visualizing this is to imagine that the concept of beer activates a beer drinking scenario, which provides slots for glasses and/or bottles of beer, but not for beer-filled pails, shoes, or lakes. In Nunberg’s terminology, our world knowledge establishes a ‘salient function’ from the ‘liquid’ sense of *beer* to a very small set of plausible count senses. This is the general mechanism underlying polysemy.

Let us now return to the denials in (3b) and (4b). My proposal is that in the context of these denials the scalar expressions *many* and *like* have bilateral meanings, and that these are instances of polysemy. This is to say, I accept the neo-Gricean premise that the lexical meaning of *many* and *like* is unilateral, but maintain that in these contexts both words have shifted into a bilateral meaning. Furthermore I maintain that this bilateral meaning is not a conventional one, and that, therefore, (3b) and (4b) are more similar to the instances of context dependent polysemy illustrated in (7) than conventionalized polysemy in (5) and (6).

It will be obvious that this analysis gives the right readings for (3b) and (4b). If *many* is interpreted in such a way that it entails *not all*, then (3b) is construed, in effect, as ‘It is not true that Fred won many but not all of the votes: he won them all’ — which seems correct. The same holds, *mutatis mutandis*, for (4b). Moreover, this proposal is based upon an interpretative mechanism that is motivated on independent grounds. What is more, polysemy is a ubiquitous phenomenon, and I see no reason to assume that it does not apply to scalar expressions. On the contrary, scalars are among the more plausible candidates for a polysemy analysis. If the lexical meaning of a scalar expression is unilateral, there will be a salient functional link to the corresponding bilateral meaning, and therefore it is only to be expected that the bilateral meaning will be available on at least some occasions. In short, what I propose is that the unilateral and bilateral meanings of the scalar adjective *warm*, say, stand to each other as do the liquid and count senses of *beer*, or, more precisely, as the ‘chicken’ and ‘chicken liver’ senses of *chicken* stand to each other.

The standard argument in favour of a neo-Gricean analysis of scalar expressions is that it is more parsimonious than an account which assumes that scalar expressions are ambiguous between unilateral and bilateral meanings. We should be wary however that considerations of lexical economy are not given too much weight. Implicature theorists like to emphasize that senses should not be multiplied beyond necessity, and although I have no objections against this slogan, lexical meanings are not on a par with scientific laws or natural kinds. If they exist in any way at all, they are individual entities, and a scientific theory does not become more complex if we double or triple the number of individuals in its ontology. There are no *methodological* reasons for preferring an astronomical theory which says that the number of planets is five to one that says that it is ten. Senses should not be multiplied beyond necessity, but postulating two senses instead of one should not be compared, if only tacitly, with, say, proliferating levels of syntactic description.

These remarks are by the way, however, because I am not committed to the claim that scalar expressions are ambiguous. I do hold that scalars are polysemous, and that, consequently, different tokens of a scalar expression may have different meanings. But ambiguity and polysemy are entirely different things, and given that *most* of the words of any natural language are polysemous, there is no reason to expect that scalars will be an exception.

The argument that I have presented in this section stands or falls with the analysis of scalar denials that I have proposed. There is, as far as I know, only one serious alternative to this

proposal, which is the unitarian account of denials advocated by Horn (1989) and van der Sandt (1991), among others.⁴ I call this a unitarian account because its crucial premise is that there is a mega-mechanism of denial (or, for Horn, metalinguistic negation) which applies not only to scalar implicatures, but to all sorts of implicatures, presuppositions — in brief, to everything that can be conveyed by means of an utterance. Methodologically speaking, it would be nice if such a mechanism could be shown to exist, because it would cover an extremely broad range of data, and it would give us an independently motivated account of scalar denial within a strictly neo-Gricean framework. However, in Geurts (in press) I argue at length that such a mega-notion of denial is untenable, because the empirical domain that it is supposed to cover is not as uniform as the unitarians claim it to be. There are several species of denial, which require different mechanisms for their explanation — and in the case of scalar denials the mechanism is polysemy.

The lexical meaning of scalar expressions

I have argued that in denials scalar expressions can have bilateral meanings.⁵ However, a denial sets up a rather special context, and I don't want to conclude on the basis of the foregoing that bilateral meanings are always available. True, it is often taken for granted that scalar expressions typically license bilateral interpretations, and it is tempting to account for this regularity on the assumption that in such cases scalars have bilateral meanings. Nevertheless, I intend to resist the temptation, and in this section I will explain why.

To begin with, the neat dichotomy between unilateral and bilateral interpretations is really too neat.⁶ Take (8) for example:

(8) Fred is bright.

In some cases, a speaker who utters this sentence will thereby signal that he is not prepared to commit himself to a stronger claim, e.g. that Fred is brilliant. Why is he not prepared to do so? One good reason is that he takes this proposition to be false. If this is what the hearer infers, then he takes (8) to implicate that Fred isn't brilliant, and thus arrives at what I have called the bilateral interpretation of *bright*. But there may be other reasons why the speaker doesn't want to commit himself to a stronger statement. For instance, he may be simply unwilling to be more specific, or he may be under some sort of obligation not to dispense too much information, or whatever. In principle, this list of possible motives can be extended *ad libitum*, and given the

⁴ See Geurts (in press) for more extensive discussion. When I say that there is only one serious alternative, I mean, of course, for the purposes of this paper. For instance, there are considerable differences between Carston's (1996) theory of denial and mine, but they don't matter here.

⁵ Later on I will show that there are more contexts in which scalars take on bilateral meanings. But even if that is the case, it doesn't affect the point I am making here, which is that scalars don't show a *systematic* preference for bilateral interpretations.

⁶ In the following I am indebted to observations by Levinson (1983), Hirschberg (1985), and Carston (1995, in press).

right circumstances any of these may be conveyed by an utterance of (8). At any rate, whether or not this list is infinitely long, the main point is that the unilateral/bilateral contrast does not exhaust the range of possible interpretations that scalar expressions may give rise to. Nonetheless, I believe that the unilateral/bilateral dichotomy has a special status, because scalars can only have unilateral and bilateral *meanings*. Consider, for example, the negation of (8), i.e.

(9) Fred is not bright.

On the default descriptive interpretation of (9), *bright* has a unilateral meaning, while on its denial interpretation it has a bilateral meaning. But although (8) may convey that the speaker doesn't know if Fred is brilliant, it seems impossible to use (9) to deny this piece of information, as the incoherence of the following exchange shows:

(10) A: Fred is bright.
 B: Fred isn't BRIGHT: you do believe that he is BRILLIANT.

If this observation is correct, how can we account for it? As follows. Suppose that the lexicalized meaning of *bright* is a unilateral one; let us label this meaning 'uni-bright'. Then the bilateral meaning of *bright* is (11a). But if (10) were a felicitous exchange, then *bright* would have to have something like the meaning in (11b).

(11) a. $\lambda x(\text{uni-bright}(x) \ \& \ \neg\text{brilliant}(x))$
 b. $\lambda x(\text{uni-bright}(x) \ \& \ \text{the speaker believes that } \neg\text{brilliant}(x))$

My suggestion is that (11a) is a natural property and (11b) is not. To be sure, I am not able to define what a natural property is, but the intuitive idea should be clear enough. Consider Goodman's (1955) famous predicate *grue*. By Goodman's definition something is *grue* iff it is observed before a certain date and is green, or is not observed before that date and is blue. Being green is a natural property, and so is being blue, but being *grue* is not. It is a property, of course, but it would never have become the meaning of *grue* except by way of definition. This is not because it is too complex. For instance, being believed to be green is a comparatively simple concept, but it is not a natural property. In general, it seems that all properties with an epistemic prefix, like $\lambda x(\text{Wilma is hoping that } x \text{ will ask her to dance})$ or $\lambda x(\text{Betty believes that } x \text{'s feet are too big})$, are non-natural properties. So, (11b) could not be a meaning of any word, because it is not a natural property. The same holds for the whole range of interpretations that scalar expressions may license, with two exceptions only: the unilateral and bilateral ones.⁷

⁷ There is a delicate (and perhaps minor) complication concerning the exact form of implicatures and meanings. Let us confine our attention to two possible implicatures that (8) may license, i.e.

- (i) It is not the case that the speaker believes that Fred is brilliant.
- (ii) The speaker believes that Fred is not brilliant.

If $\lambda x(\text{the speaker believes that } \neg\text{brilliant}(x))$ is a non-natural property, as I suggest, then strictly speaking the implicature in (ii) can never become part of the meaning of *bright*. So

But let me return to my main line of argument, which is about the hypothesis that bilateral interpretations might generally (i.e. not only in denials and perhaps a few other cases) be traced back to bilateral meanings. It is widely assumed that scalars typically license bilateral interpretations. This is why neo-Griceans insist that scalar implicatures are *generalized* conversational implicatures. But is it true? We have observed that, in addition to the unilateral and bilateral interpretations, there is a whole range of possible interpretations that scalars may give rise to. Is there any empirical evidence for the claim that it is the bilateral interpretation of (8) that is the most likely to occur rather than, say, the interpretation on which the speaker simply has no opinion as to whether Fred is not just bright but actually brilliant?

According to the neo-Gricean theory, the lexical meaning of *bright* is unilateral, and thus a speaker who utters (8) merely says that Fred is uni-bright, i.e. his statement does not entail that Fred is not brilliant. Now if we want to account for any of the non-unilateral interpretations observed above, the familiar Gricean train of reasoning will have to be set in motion. But regardless what it is we want to derive, we need to assume that it is relevant (to the purpose of the discourse) that the hearer knows if Fred is brilliant. For quantity implicatures, and indeed conversational implicatures in general, are conditioned by considerations of relevance, as Grice himself already hinted in his formulation of the first quantity maxim,⁸ and which has since been emphasized by many authors, especially representatives of relevance theory. So if we want to show that generally speaking *bright* licenses a scalar implicature, it should be the case that generally speaking it is relevant to know if someone who is said to be bright is in fact brilliant. As far as I can tell there is no justification for this assumption.

There is, admittedly, a strong temptation to assume that (8) generally will give rise to construals that convey more than a unilateral interpretation does. But that temptation arises because of the way implicature theorists have asked their questions. For they have asked themselves questions like: ‘Would an utterance of (8) imply that Fred is brilliant?’ But to ask this question is to *presuppose* that (8) is uttered in a context in which it is relevant to know if Fred is brilliant or not. And if we restrict our query to this type of context it is hardly surprising that the unilateral interpretation should turn out to be too weak in the general case.

The upshot of these observations is that we simply do not know if scalar expressions license bilateral interpretations on a regular basis. So the situation we find ourselves in is the following. It may safely be assumed that in the vast majority of possible contexts in which (8) is felicitously uttered, it will be construed as entailing that Fred is uni-bright; this much is accounted for by the hypothesis that *bright* actually means ‘uni-bright’. Furthermore, it seems plausible that in a large subset of these contexts (8) will convey some extra information, for instance, that the speaker doesn’t know if Fred is brilliant, that it would be against regulations to tell the hearer if Fred is brilliant, and so on. But we don’t know how large this subset

we cannot say, strictly speaking, that some implicatures can and others cannot become part of the meaning of a scalar expression. What we might say instead is that in some cases, though not in all, a scalar can entail what otherwise is implicated. Thus, if $\lambda x(\neg\text{brilliant}(x))$ were part of the meaning of *bright*, (8) would entail that Fred is not brilliant, and an utterance of this sentence would imply (ii). The point that I want to make is then that there is no natural property P such that, if P were part of the meaning of *bright*, an utterance of (8) would imply (i) because of this meaning.

⁸ ‘Make your contribution as informative as is required (for the current purposes of the exchange).’ Grice (1989: 26)

actually is, i.e. how often any one of these additional pieces of information is conveyed, and we don't know how often a bilateral interpretation is conveyed.

I take it that meanings will only be lexicalized if it is cost-effective. Although the word *chicken* may be used to express the concept 'chicken liver', as Nunberg's example (7c) demonstrates, this use is so rare that 'chicken liver' is not among the lexicalized meanings of *chicken*. In the foregoing I have argued that there is no reason to assume that it is any different with the bilateral interpretations of scalar expressions. We have no more reason to assume that these are lexicalized than we have for assuming that 'chicken liver' is a lexical meaning of *chicken*. By the same token, there is no evidence for the neo-Gricean hypothesis that scalars trigger bilateral interpretations through generalized conversational implicatures. I will have more to say about this, but first let me make a few remarks about

The cardinal issue

The picture I have outlined in the foregoing deviates from the orthodox neo-Gricean view in two respects. I agree with the neo-Griceans that, in general, the lexical meaning of a scalar expression will be unilateral, and that bilateral interpretations are to be explained in terms of conversational implicature, but on the other hand I hold, first, that on some occasions scalars have bilateral meanings, and secondly, I am not convinced that scalar implicatures are generalized conversational implicatures. In this section I want to make a final amendment to this picture, to the effect that cardinals (*one, two, ...*) are not scalar expressions.⁹ If this position, which I believe has become the majority view, is correct, then it is of broader interest, because for a long time cardinals have been routinely used to illustrate all sorts of claims about scalars in general.¹⁰

The literature on this subject yields quite a number of arguments for taking a separatist stance on the cardinal issue. Some of these are not very effective, in my opinion, and I will confine my attention to what I take to be the most persuasive one, which is based on the observation that scalars have a peculiar distribution pattern and combine with modifiers that, presumably, require them to have bilateral meanings.

$$(12) \left\{ \begin{array}{l} \text{at least} \\ \text{at most} \\ \text{more than} \\ \text{less than} \\ \text{exactly} \\ \text{roughly} \end{array} \right\} \left\{ \begin{array}{l} \text{two hundred} \\ \text{half of the} \\ \text{?most (of the)} \\ \text{?few (of the)} \end{array} \right\} \text{ ships}$$

The pattern is clear: there are several modifiers that only combine with cardinals and related expressions like 'half of', and that do not combine with scalars like *few, most*, and so on

⁹ Alternatively I might have claimed, of course, that there are substantial differences between cardinals and other scalars. But this is just a terminological matter.

¹⁰ For discussion, see Sadock (1984), Löbner (1985), Koenig (1991), Horn (1992), Carston (in press). To the best of my knowledge, the only recent defense of the position that cardinals should be treated on a par with scalars is Levinson (1997).

(these sweeping claims will be qualified below). Furthermore, it is intuitively clear that the argument of, say, *at least* should be a precise value, and is thus required to have a bilateral meaning.

I think this is a very good argument indeed, but it should be noted that it is only about 98% watertight. This is slightly less good than, for instance, Koenig (1991) takes it to be, who maintains that if *six* has a unilateral meaning, we ‘cannot make sense’ of an expression like *at least six sheep*. This overstates the case, if only marginally. The point is not that there is no way we could make sense of such expressions, but rather that they would force us to pervert the semantics of *at least* and its congeners. To see this, let us picture the meaning of *six* as being either the number 6 (the bilateral view) or the series of natural numbers beginning with 6, i.e. [6, 7, 8, ...] (the unilateral view). Now consider the semantics of *at least* and *at most*. On the bilateral view, the interpretation of ‘at least six’ is obtained by mapping 6 onto [6, 7, 8, ...], while *at most* maps 6 onto [... , 4, 5, 6]. On this view, *at least* and *at most* are mirror images, which is what one should expect. On the unilateral view, however, these modifiers are quite different from each other, *at least* mapping [6, 7, 8, ...] onto itself, and *at most* mapping the same sequence onto [... , 4, 5, 6]. This seems wrong for two reasons: it is implausible that *at least* should be semantically empty, and there is a strong intuition that *at least* and *at most* should only differ in orientation, not in type. Or consider the semantics of *approximately*. On the bilateral view, *approximately six* is interpreted by mapping 6 onto a vague area around itself, something like this: [... , 4, 5, 6, 7, 8, ...]. On the unilateral view, by contrast, the sequence [6, 7, 8, ...] is mapped onto a vague area around its first term. This is not impossible, but it does seem quite unlikely.

On the bilateral analysis of cardinals that I endorse, modifiers like *exactly*, *precisely*, and so on, are identity functions, semantically speaking: they map a number onto itself. I don’t think this is an counter-intuitive consequence. I find it rather plausible that these modifiers merely serve to signal that their arguments are to be construed strictly. In this respect, these modifiers are rather like the floating universal quantifier in,

(13) They all became quite ill.

The truth conditions of this sentence would not be affected if we excised *all*. But the resulting sentence would be easier to construe loosely than (13). *Exactly* seems to perform a similar duty, as other uses confirm: ‘exactly in the middle’, ‘exactly the same’, and so on. Finally, it bears emphasizing that whereas the bilateral analysis of the cardinals entails that *exactly* is semantically empty, the unilateral analysis entails that this holds for *at least*. Faced with this dilemma, the choice should be easy.

The foregoing argument uses expressions of the form ‘at least X’, ‘exactly X’, ‘less than X’, and so on, to show that X must have a bilateral meaning. But as one of my premises is that expressions may acquire new meanings in a given context, this does not yet prove that X’s lexical meaning must be bilateral. It might be argued in principle that the lexical meaning of cardinals is unilateral, and shifts into a bilateral meaning in the context of *at least*, *exactly*, and so on. And this position would have some *prima facie* plausibility in view of examples like the following:

- (14) a. Dinah caught at least some mice.
b. He lost more than a few dollars.

- c. The water was more than warm: it was hot.

However, although scalars do combine with some of these modifying expressions, they are not nearly as free in this respect as cardinals are. Compare:

- (15) a. Dinah caught at least {?few / ?many / three} mice.
b. Dinah caught at most {?some / three} mice.
c. He lost less than {?a few / thousand} dollars.
d. He lost exactly {?a few / thousand} dollars.

This array of facts does not lend itself to a straightforward explanation, as far as I can see, but the crucial datum is that, whereas cardinals combine quite freely with *exactly* and its ilk, scalars don't. Therefore, it is plausible to assume that in the case of the cardinals bilateral meanings are actually lexicalized, while in scalars they are not.

On the notion of generalized conversational implicature

In the foregoing I have argued that there is no compelling evidence for the claim that scalar expressions have a lexicalized bilateral meaning. But the considerations that lead to this conclusion apply more generally: there is no evidence that, in general, the bilateral interpretation that a scalar expression α may license is caused by any type of content conventionally associated with α . In particular, there is no reason to believe that scalar implicatures are generalized conversational implicatures, as claimed by Levinson, Gazdar, Horn, and others. GCIs are defaults. If it is claimed that an expression α licenses a GCI to the effect that ϕ is true, then α will trigger the inference that (or convey that) ϕ is true *unless* there is evidence to the contrary. Thus, if there is a 'not all' GCI associated with *some*, then any utterance of a sentence containing *some* will imply that the corresponding sentence with *all* is not true, unless there is evidence indicating that this implication is not intended. It follows from this, evidently, that GCIs are conventional(ized) elements of interpretation, although in this case conventional does not entail arbitrary, since GCIs are always motivated by pragmatic considerations. In the following I will consider some of the arguments that neo-Griceans have adduced in defense of the notion of GCI. In particular, I want to consider the arguments for claiming that scalar implicatures are GCIs, and argue that these are not very convincing. However, as scalar implicatures are by far the most likely candidates for the status of GCI, I believe that my argument casts doubt on the very notion of GCI.

In a sense, the strongest version of the claim that there are GCIs is that there is a projection problem for implicatures just as there is a projection problem for presuppositions. As Levinson puts it:

In short, we have a projection problem: we need to be able to state the conditions under which a potential implicature generated from a part of a sentence is in fact inherited by the whole. (Levinson 1997, chapter 2: 93)

This claim, I will argue, is misleading if not downright false.¹¹

What does it mean to say that there is a projection problem for presuppositions?

Presuppositions are triggered by certain lexical elements and syntactic constructions (definite articles, factives, clefts, and so on), and they exhibit projection behaviour in the sense that have a strong tendency to ‘escape from’ embedded positions. In general, a presupposition holds regardless the level of embedding at which it is triggered, as the following examples show:

- (16) a. $\left. \begin{array}{l} \text{If he has seen the papers, then} \\ \text{Let's not exclude the possibility that} \\ \text{Wilma refuses to believe that} \end{array} \right\} \text{the PM knows that the people are hungry.}$
b. The people are hungry.

The factive verb *know* triggers the presupposition that its complement is true, and as (16) illustrates this presupposition tends to go through regardless: any of the three variants of (16a) will normally be taken to imply that (16b) is true. There is a tendency in the literature to equate presuppositions with those inferences that tend to escape from the scope of a negation operator, but this is only part of much larger pattern, since they tend to escape from *any* embedded position, even if it is created by an attitude verb. This should be borne in mind when one is trying to determine whether or not a certain inference is of a presuppositional nature, because the negation test doesn't work if a presuppositional expression is a positive polarity item, for instance (see van der Sandt 1988 for discussion).

This is one half of what it means to say that presuppositions show projection behaviour: they tend to hold no matter where they are triggered. The other half is that this general rule does not apply under certain, rather special circumstances. For example:

- (17) If the people are hungry, then the PM knows that the people are hungry.

Here the presupposition triggered by the factive verb is blocked in the antecedent of the conditional, as a result of which the sentence as a whole does not imply that the presupposition is true.

So to say that an inference displays projection behaviour is to say that it will normally go through regardless the level of embedding at which it is triggered, unless special circumstances obtain, as is the case in (17).

It is always tricky to define linguistic phenomena in a theory-independent manner, but in the case of presuppositions we can at least give a passable working definition: presuppositions are precisely those inferences that show projection behaviour. At the very least one can say that, if it can be shown that a certain inference shows projection behaviour, we have rather strong evidence that some form of presupposition is involved. And let me add that, as linguistic notions go, the pre-theoretic notion of presupposition is comparatively unproblematic.

Since a number of years there is also a clear majority view on the approximate direction that a theory of presupposition should take. I will not attempt to define this view, because there is only one aspect of it that is relevant to my current purposes. It is that it is widely accepted that presuppositions can never be *cancelled*. If a presupposition seems to disappear, as in (17), it is actually bound or satisfied or accommodated below sentence level, but presuppositions are never cancelled (as was sometimes assumed in the earlier literature on this subject, e.g. Gazdar

¹¹ I have made this point before (Geurts 1995). In the following I expand on that discussion.

1979). One indication that this assumption is correct is that presuppositions triggered in simple sentences and in other entailed positions can never be made to disappear, but since this assumption is not really a controversial one, I will not attempt to defend it here.

Let us now return to the notion that there is a class of implicatures that show projection behaviour. If this is to be taken literally, then we have strong evidence that these implicatures are presuppositions, for, as we have seen, projection behaviour is the hallmark of presuppositions. It is evident, however, that this is not what neo-Griceans want to claim. Apparently, therefore, when it is said that these implicatures show projection behaviour, the intended claim is weaker, say, that there is a significant similarity between these implicatures and presuppositions. But even this watered-down version is incorrect. To begin with, there are a number of cases where we seem to observe projection behaviour of implicatures, but which are easily accounted for in different terms. For example:

(18) {Barney believes that / Perhaps} the water is warm.

Suppose, if only for the sake of the argument, that an utterance of (18a) or (18b) would normally convey that, according to the speaker, the water isn't hot. Would this support the claim that the scalar implicature associated with *warm* shows projection behaviour? No, because this observations is straightforwardly explained without this assumption. If a speaker utters one of these sentences, he implies that he doesn't know if the water is warm (this is what Gazdar calls a clausal implicature). But then the speaker doesn't know if the water is hot either. This is not quite the same as saying that, according to the speaker, the water isn't hot. But this is a minor hurdle, which can be taken with the familiar Gricean apparatus of pragmatic inference. So the observation, correct or not, can be accounted for in strictly Gricean terms without assuming that there is a GCI associated with *warm*.

But I suspect that our initial observation was not correct, at least not in general. Cf.

(19) $\left\{ \begin{array}{l} \text{Barney suspects that} \\ \text{Barney doubts that} \\ \text{It is extremely likely that} \\ \text{Can you tell me if} \\ \text{Let's go and see if} \end{array} \right\}$ the water is warm.

According to my intuitions, none of these sentences definitely imply that the water isn't hot, and at least some of them definitely don't imply this. By contrast, each of these sentences clearly implies, in the absence of evidence to the contrary, that there is water (according to the speaker). Even if our intuitions about the exact interpretation of the scalar *warm* are somewhat insecure, our intuitions about the presuppositional expression *the water* are crystal clear. Which confirms that the evidence that implicatures show projection behaviour is tenuous at best.

If it cannot be shown that conversational implicatures have a tendency to escape from embedded positions, as presuppositions have, then *a fortiori* it cannot be shown that they are literally blocked or suspended under some circumstances. Cases like the following, for example, do not show that some implicatures exhibit projection behaviour:

(20) a. The water is warm, if not hot.
 b. The water is warm, and perhaps even hot.

- c. The water is warm. In fact, it is hot.

If we want to use examples like these to argue that the scalar inference associated with *warm* shows projection behaviour, then it first has to be demonstrated that there is strong tendency for this inference to be triggered and projected, and I have argued that this is by no means clear. Whenever such examples are described in terms of blocking and/or suspension, it is *presupposed* that scalar inferences show projection behaviour; they don't provide independent evidence for this claim. Indeed, examples like (20b, c) provide additional evidence *against* the hypothesis that scalar implicatures show projection behaviour. I pointed out above that presuppositions triggered in simple sentences and in other entailed positions can never be 'cancelled'. For instance, the following are simply contradictory:

- (21) a. ?The water is warm, but there is no water.
 b. ?The water is warm. In fact, there is no water.

However, neo-Griceans claim that in the context of (20b, c) the implicature that the water is not hot is cancelled. If that is the case, this is yet another point at which the analogy between presuppositions and implicatures breaks down.

The upshot of these arguments is that there is no evidence that implicatures are like presuppositions in that they show projection behaviour. But actually I have already begun to argue against the weaker claim that scalar implicatures are GCIs. For one of the main arguments in favour of this position is precisely that implicatures can be blocked or suspended. If an inference is blocked or suspended it must have been triggered first, otherwise there wouldn't have been anything to block or suspend. However, there is no independent evidence that in 'suspension' contexts like (20a) anything like a scalar implicature is triggered at all, and even if we should grant that in 'blocking' contexts like (20b, c) a scalar implicature is first derived and then cancelled, it is not clear why we should say that this is a GCI. Blocking and suspending contexts can perhaps be accounted for in terms of GCI, they provide no independent evidence for this notion.

Another argument adduced by Levinson (1997) is based upon certain patterns of lexicalization documented in detail by Horn (1989). If we take the positive duals *all* and *some*, we observe that the (external) negation of the latter is lexicalized while that of the former is not. The same holds for other pairs of duals, like *necessary/possible*, *and/or*, and so on. In all cases, there is a lexical gap in the same place, as the following diagram shows:

<i>all</i> <i>necessary</i> <i>obligatory</i> <i>or</i>	<i>none</i> <i>impossible</i> <i>forbidden</i> <i>neither/nor</i>
<i>some</i> <i>possible</i> <i>allowed</i> <i>either/or</i>	<i>not all</i> <i>n. necessary</i> <i>n. obligatory</i> <i>not both</i>

Horn's explanation for this pattern, which is found not only in English but cross-linguistically, is that *some*, *possible*, etc. tend to implicate *not all*, *not necessary*, etc., which obviates the

need for special lexical entries for the latter group. Levinson endorses this analysis, and maintains that it supports the notion of GCI:

Anyone who rejects the notion of a GCI or a preferred interpretation will have to have another explanation for this lexicalization pattern. The nonce-implicature theorist cannot explain it — according to any such theory, there are no general tendencies to be found, or if there are they have the status of mere behavioural tendencies, playing no role in the systematic generation of implicatures. (Levinson 1997, chapter 1: 67)

This is incorrect, however, for the crucial ingredient in Horn's analysis is not that *some* (and likewise *possible*, *or*, and so on) comes with a GCI to the effect that *all* does not hold. The keystone of Horn's explanation is that, generally speaking, when a speaker should want to use *not all* he can use the shorter form *some* in its stead, because on such occasions *some* will implicate 'not all'. This, together with the assumption that positive concepts have a stronger tendency to be lexicalized than negative ones, suffices to explain the lexicalization pattern observed by Horn. In order to explain this pattern, it only need to be assumed that, in general, 'not all' can be conveyed by means of *some*; what Levinson needs to prove however is that in general *some* conveys 'not all' and that moreover this information has become conventionally associated with the word *some*. Therefore, this argument is simply fallacious.

In this section I have criticized the idea that scalar implicatures are GCIs, not on the grounds that it is inconsistent or because it makes the wrong predictions, but because I can find no compelling evidence to support it. If this isn't my fault — that is, if there really is no strong evidence for this hypothesis — then the notion of GCI is in a rather precarious condition, for if there is any variety of implicature that epitomizes the notion of GCI it is the scalar implicatures. Although it goes beyond the thematic confines of this paper, I do want to take this argument a little bit further. Levinson distinguishes between three types of GCI, which he labels Q (quantity), M (manner), and I (informativeness) (a similar classification is put forward by Horn 1989, for example). Scalar implicatures, which exploit Grice's first quantity maxim, belong to the first class. I-implicatures exploit the second maxim of quantity. The following examples (from Levinson 1997, chapter 1: 29) belong to this group:

- (22) a. John's book is good.
I-implicature: the book John read, wrote, borrowed, as appropriate.
- b. The picnic was awful. The beer was warm.
I-implicature: the beer was part of the picnic.
- c. John and Jenny bought a piano.
I-implicature: they bought it together.

Let us grant that these are conversational implicatures (I find it hard to accept that they are, but that is an issue I don't want to go into here). But are they *generalized* conversational implicatures? Take the NP *John's book*. We may suppose that the grammar of English underdetermines the interpretation of this construction on any given occasion. Perhaps its meaning is something like: 'Here is John and there is a book and there is some sort of relation between the two'. This is not very specific, to put it mildly, and therefore contextual knowledge and world knowledge must be called upon to figure out whether John read, wrote, or borrowed

the book. Everybody will agree that something like this story is probably true. But where does the GCI come in? Nowhere, as far as I can tell. Similar remarks apply in the case of (22b). (22c) is a rather different case, but here too I fail to see what the notion of GCI is supposed to explain, even if it is granted that the preferred interpretation of this sentence is that John and Jenny bought a piano together. And I don't see why this should be granted in the first place. I will have this preference if I am told, say, that John and Jenny share a common household, but not if they live on different continents.

Levinson is aware that there might be thought to be a problem here:

In what sense are these I-inferences *generalized*? Most of these inferences interact with shared background assumptions, which might in principle vary, and thus the inferences might have none of the cross-context, even cross-linguistic, invariance that are the hallmarks of GCIs. But at a sufficient level of abstraction, it is quite clear that the kinds of inferences here collected [...] do hold as preferred interpretations across contexts, and indeed across languages. (Levinson 1997, chapter 2: 51)

But this will not do, because if we go to a sufficient level of abstraction, these inferences will be abstracted from, too. Consider *John's book* again. In some contexts this will be interpreted as 'the book John owns', in others, as 'the book John wrote', and so on. But if we abstract from the particulars of these contexts, we perforce abstract from these inferences, too.¹² Similarly, the bridging inference in (22b) hinges upon assumptions that hold in some contexts but not in others. It is easy to imagine cultures or subcultures in which the concepts of picnic and beer have nothing to do with each other, or are related in ways that deviate from what Levinson, apparently, takes to be the standard. If we leave such context-specific assumptions out of account, the bridging inference simply evaporates.

The notion of generalized conversational implicature is a dubious one, not only as applied to scalars, but across the board.

References

Carston, R. 1995: Quantity maxims and generalised implicature. *Lingua* 96: 213-244.

¹² If there is a default interpretation of the genitive construction, then it surely is that it denotes possession. This is rather like the case of *secretary*, which by default will be interpreted as 'female secretary'. Levinson suggests that this is a GCI, too, but again I fail to see on what grounds. I do admit that in the absence of evidence to the contrary I tend to assume that secretaries are women, and it may be that most people employ this default rule. However, although GCIs are default rules, not all default rules are GCIs, and I don't see what is gained by claiming that, in addition to the default rule that secretaries are women, there is a GCI associated with *secretary* to the effect that this *word* implies 'female secretary'. I have the impression that Levinson's (and not only Levinson's) use of the concept of GCI is at least sometimes marred by an equivocation between this notion and the notion of default inference *tout court*.

- Carston, R. 1996: Metalinguistic negation and echoic use. *Journal of Pragmatics* 25: 309-330.
- Carston, R. in press: Informativeness, relevance and scalar implicature. In: R. Carston and S. Uchida (eds.), *Relevance Theory: Applications and Implications*. John Benjamins, Amsterdam.
- Gazdar, G. 1979: *Pragmatics: Implicature, Presupposition, and Logical Form*. Academic Press, New York.
- Geurts, B. 1995: *Presupposing*. Doctoral dissertation, University of Stuttgart.
- Geurts, B. in press: The mechanisms of denial. In: *Language*.
- Geurts, B. and R.A. van der Sandt, in press: Domain restriction. In: P. Bosch and R.A. van der Sandt (eds.), *Focus: Linguistic, Cognitive, and Computational Perspectives*. Cambridge University Press.
- Goodman, N. 1955: *Fact, Fiction, and Forecast*. Harvard University Press, Cambridge, Mass.
- Grice, H.P. 1989: *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass.
- Hirschberg, J. 1985: *A Theory of Scalar Implicature*. Ph.D. thesis, University of Pennsylvania.
- Horn, L.R. 1989: *A Natural History of Negation*. Chicago University Press.
- Horn, L.R. 1992: The said and the unsaid. In: C. Barker and D. Dowty (eds.), *Proceedings of SALT II*. Ohio State University. Pp. 163-192.
- Kempson, R.M. 1986: Ambiguity and the semantics-pragmatics distinction. In: C. Travis (ed.), *Meaning and Interpretation*. Blackwell, Oxford. Pp. 77-103.
- Koenig, J.-P. 1991: Scalar predicates and negation: Punctual semantics and interval interpretations. *Papers from the 27th Regional Meeting of the Chicago Linguistic Society, Part 2: The Parasession on Negation*. Chicago Linguistic Society. Pp. 331-344.
- Levinson, S.C. 1983: *Pragmatics*. Cambridge University Press.
- Levinson, S.C. 1997: *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Ms, Max Planck Institute for Psycholinguistics, Nijmegen.
- Löbner, S. 1985: Drei ist drei: zur Bedeutung der Zahlwörter. In: W. Kürschner and R. Vogt (eds.), *Akten des 19. linguistischen Kolloquiums*. Niemeyer, Tübingen.
- Nunberg, G.D. 1978: *The Pragmatics of Reference*. Indiana University Linguistics Club.
- Sadock, J. 1984: Whither radical pragmatics? In: D. Schiffrin (ed.), *Meaning, Form, and Use in Context: Linguistic Applications*. Georgetown University Press, Washington. Pp. 139-149.
- Van der Sandt, R.A. 1988: *Context and Presupposition*. Routledge, London.
- Van der Sandt, R.A. 1991: Denial. *Papers from the 27th Regional Meeting of the Chicago Linguistic Society, Part 2: The Parasession on Negation*. Chicago Linguistic Society. Pp. 331-344.