

To appear in:

KLAUS VON HEUSINGER and KEN TURNER (eds.)

WHERE SEMANTICS MEETS PRAGMATICS:
THE MICHIGAN STATE UNIVERSITY PAPERS

in the series

CURRENT RESEARCH IN THE SEMANTICS/PRAGMATICS INTERFACE

<http://www.elsevier.com/locate/series/crispi>



AMSTERDAM

A ‘JUST THAT’ LEXICAL MEANING FOR *most*

Mira Ariel, Tel Aviv University, Israel

1 Introduction*

Following Horn (1972), the lexical meaning assumed by the received view for scalar quantifiers such as *most* (namely, only lower-bounded) quite routinely underdetermines the meaning it is actually assigned by interlocutors (upper-bounded as well). Linguists prefer to derive the upper bound as a conversational implicature. *Most* is said to code ‘more than half’, but since the speaker simultaneously generates a ‘not all’ implicature, the conveyed meaning is both lower and upper-bounded (‘more than half, but not all’). The pragmatic inference which provides the upper bound must be “default”, because it is commonly needed for interpreting *most*. A pragmatic inference is then required in order to bridge the gap between the lexical meaning (‘more than half’) and the conveyed meaning of *most* (‘more than half, but not all’). While Grice’s 1975 work has made it obvious that conveyed meanings are often mediated by pragmatic enrichments of various sorts, it should be equally obvious that we should avoid such solutions for lexical items as much as possible. Such gaps should be assumed as a last resort, preferably only for cases involving ad hoc pragmatic inferences (Particularized Conversational Implicatures), and not for pragmatic inferences which seem almost obligatory (Generalized Conversational Implicatures). Being routinely generated, the pragmatic enrichment associated with *most* (‘not all’) is of the latter type. I argue below that we can account for all the interpretations of *most*, including ones the received view is incapable of accounting for, without assuming a gap between the lexical meaning of *most* and its prevalent interpretation. On my analysis *most* means ‘a proper subset which is the largest subset, given any partitioning of the complement set (into one or more subsets)’. This is a ‘just that’ lexical analysis for *most*, namely, one where the extension denoted by its lexical meaning is just what its common conveyed meaning is: ‘More than half and less than all’.

Consider the following, typical (1) and less typical (2) examples with quantifier *most*:

- (1) **Most** (Israelis) decided for peace. Me too (Originally Hebrew, bumper sticker).

*I wish to thank Uriel Cohen, John Du Bois, Edit Doron, Laurence Horn, Nirit Kadmon, Fred Landman, Aldo Sevi and the participants of the “Where semantics meets pragmatics” workshop for fruitful discussions about *most*. I’m also grateful to my two anonymous referees, as well as to Ken Turner, for their helpful criticisms. I thank Gila Batori for the statistical analyses of the data.

- (2) MONTOYA: ... would you say that **most** of you,
 without getting personal here,
 are products=,
 ... uh=,
 in terms of birth,
 ... of the sixties?
 ... **All** of you were born in the sixties? <SBC: 012>.

Example (1) is typical for my data in that the addressor is trying to convince the addressee to adopt some stance or action based on a generalization which holds true of some majority (“majority can’t be wrong”). (2) is less typical, because *most* is more easily seen as compatible with the speaker believing that possibly all of his addressees were born in the sixties, as he indeed proceeds to indicate.

Most is associated with several different interpretations (I indicate in parentheses which examples the specific interpretation is relevant for):

- (3) a. More than half (1+2).
 b. Less than all (1).
 c. Possibly all (2).
 d. A noteworthy amount (1+2).¹
 e. (Based on d.): Some consequence/conclusion is being justified by the speaker, e.g., we should all decide for peace (1+2).

It seems quite obvious that interpretations (3d) and (3e) are pragmatically derived, whereas (3a), the lower boundary on *most*, is lexically specified. But what about (3b) and (3c)? The received view (ever since Horn 1972, see also Horn 1984, 1989, 2005) has been that (3c) is lexically allowed for: *Most* is only lower-bounded, and is therefore compatible with ‘all’ being the case, just as it is compatible with any value above half (e.g., 87%). (3b) is pragmatically derived as a scalar implicature, due to a comparison of *most* with the more informative *all* (on a conventional Horn scale), which the speaker could have used. Since she did not choose the more informative *all*, she must not intend ‘all’.

Other lexical-pragmatic divisions of labor are also conceivable. Hamilton (1860) (as cited in Horn 1989:207) proposed an ambiguity analysis for *some*. We could propose a similar analysis

¹To see that *most* tends to denote a noteworthy amount, observe the objection (in b) to the use of *most* (in a), just because the majority is rather small (http://www.talkleft.com/new_archives/004742.html, 12.20.2003):

- (i) a. An ABC News poll shows that **most** people support a life over death sentence for John Lee Malvo, who was a juvenile at the time of the sniper attacks.
 Given a choice, 52 percent say they’d prefer a sentence of life in prison with no chance of parole for Malvo, while 37 percent prefer the death penalty. (http://abcnews.go.com/sections/us/WorldNewsTonight/malvo_death_penalty_poll_031219.html)
 b. Fifty two percent is **not** “**most**”.
 (http://www.talkleft.com/new_archives/004742.html, 12.21.2003).

Presumably, the poster of (b) would find *more than half* a more appropriate description for 52%. The original use, I assume, views 52% as noteworthy, because the opposed view is only endorsed by 37% , rather than by 48% of the people surveyed.

for *most*, so that in addition to a lower-bounded meaning, it also has a bilateral meaning. For the bilateral meaning, both (3a) and (3b) are lexical, thus ruling out interpretation (3c) altogether. Since even (1), (where *most* is interpreted as upper-bounded), is not false should 'all' be the case, i.e., should all Israelis have decided for peace, we can safely dismiss the bilateral account.

I too have recently argued for a view which is (almost) the opposite of the received view, namely, that something like (3b) is lexically specified and that (3c) is mediated by pragmatic inference (see Ariel 2003, 2004). Based on 127 corpus examples (all the quantifier *most* examples in the Santa Barbara Corpus of Spoken American English, the London Lund Corpus and Morris 1994:71–192) and on questionnaire data (see also 2.2 below), I have argued that: (a) Pragmatics cannot provide the upper bound on *most*, and moreover, (b) Semantics should provide the upper bound on *most*. A 'just that' semantic analysis for *most* stipulates that its prevalent interpretation is also its lexical meaning. This article focuses on the second part of the argument, namely that semantics should and in fact can account for the upper bound of *most*, and that *most*'s compatibility with 'all' is only pragmatically enabled (but see also Ariel 2003, 2004).

A note is here called for on my repeated use of the term lexical, rather than semantic, and on how I understand the term upper bound(ed). My argument is that *most*'s lexical meaning is upper bounded, not that the semantic (i.e., truth-conditional) interpretation of *most* utterances is necessarily always upper-bounded. Following Relevance Theory (Sperber & Wilson 1986/1995), I take it that pragmatic enrichments (such as 'not all' or 'possibly all') can contribute to the truth conditions of propositions, and therefore can constitute part of the overall (referential) semantics of the proposition. Put differently, while the linguistic semantic analysis of *most* is upper bounded on my analysis, the referential semantic interpretation need not be so (see Carston 1999 for the concept of referential semantics). Here's a relevant example, where *most*'s referential semantic interpretation is 'at least most, possibly all', even though I maintain that its lexical meaning is upper bounded:

- (4) The target date for the meeting is Jan. 17 in Los Angeles, provided **most** of the Hall of Famers can make it (International Herald Tribune, 12.24–25, 2002, p. 16).

Second, as will become clearer when I discuss the compatibility of *most* with 'all', I don't see the lexical upper bound on *most* as an 'all-exclusion' interpretation. In other words, when I say that *most* denotes less than all the relevant set, it does not follow that the predicate asserted to be true for *most* is necessarily false of the part not denoted. This would, of course, be the bilateral position which I am rejecting. More on this in section 3.

Before we address the main topic of this paper, I should briefly summarize the arguments made in Ariel (2004) for my claim that pragmatics cannot account for the (routine) upper bound on *most*. 'Not all' (of whatever strength) is not a frequent implicature in the data I examined. The reason is that the presumed implicature is either not relevant, or worse, it would actually undermine the speaker's communicative intent in uttering the *most* proposition. I argued that 'not all' is relevant and communicated only when there is an expectation for 'all' to be the case. As it turns out, we don't normally have such high expectations. We don't expect to be able to generalize over all instances. *Most* is considered a very significant amount, one we can build on in justifying consequences following the *most* proposition (this is especially clear in ads stating that most people prefer some product). Note that generating the 'not all' implicature actually defeats the speaker's purpose in using the *most* proposition, since the minority 'not all' set of instances actually constitutes a counter-example to the speaker's point. Speakers usually use the

most proposition in order to support some further action/conclusion they wish to convey. In the typical (1), for example, the addressor would like to convince all Israelis to decide for peace. The majority of Israelis are the example-setters for all of us in this case. The addressor certainly does not want us to follow the example set by the minority group of Israelis who did not decide for peace (see Ariel 2004 for more examples and discourse statistics). Nonetheless, in these very cases where no implicature is generated, *most* is seen as upper-bounded. Since it relies on implicature generation, the received view, I argued, cannot account for the prototypical *most* cases.

In response to these findings, Horn (2005) insists that an implicature is generated. He proposes to shift the basis of the scalar implicature to the Maxim of Quality, in order to defend his original position that the upper bound on *most* is pragmatically derived in these cases.² On this revised view, *most* conveys ‘not all’, because otherwise, using *most* when ‘not all’ is the case violates Quality (the reason is that *most* is also compatible with ‘all’, which is here false). There are a few problems with this position. First, it means that ‘not all’ is always inferred by the addressee, regardless of its relevance (unlike the original proposal). Scalar implicatures can never be cancelled implicitly by context (as in (2), even before we hear the last sentence). Second, it is not clear to me how using *most* is a breach of Quality under the received view. If under the received view *most* denotes ‘more than half’ then it cannot be false for ‘most, but not all’ cases. The fact that it allows also for the 100% value to be true, which is false in this case (say we’re talking about 85%), is quite parallel to the fact that *most* also allows it for 75% to be true, even though it is false in the assumed case. No one would want to say that *most* violates Quality in such a case. I think that the problem is that semanticists and pragmatists have singled out the 100% value as the only relevant value worthy of consideration when scalar quantifiers are accounted for. A clear Gricean case of Quality violation are ironies, where what the speaker asserts is blatantly false, in which case the addressee reasons that the speaker actually wants to convey something different and truthful. A speaker using *most* when ‘not all’ is the case cannot, however, be seen as saying anything false, and hence, I don’t see how Quality has been breached by such a speaker. The use of *most* can only be seen as less than maximally informative on the Neo-Gricean account, and hence, must generate a Quantity implicature, which is cancelable by Relevance. This has been the original Neo-Gricean position (see also Levinson 2000), and for a good reason.

Third, even if my two arguments above fail, all things being equal, we should avoid the assumption of a conversational implicature (Generalized conversational implicatures included) which must routinely apply in violation of Relevance (the counts reported on in Ariel 2004 are that for at least 74% of the *mosts* in the data the implicature would have to be generated despite its irrelevance and counter-productive effect in the conversation). Even if both the received view and my circumbounded view can account for all the data, it seems preferable to involve pragmatic inferencing in the minority of the cases (at most 26%), rather than in the majority of the cases.³

²Horn argues that we sometimes implicate information against our best interest, but I have argued in Ariel (2004) that those cases are interactionally quite different from the unmarked *most* case. The “reluctant” implicatures are all cases where the more informative term (*all* in our case) is expected. This is precisely not the case in the unmarked *most* utterance. See Ariel (2004) for details.

³While agreeing with me that we should reserve implicatures to the minority rather than to the majority cases, one of my referees counters that it is possible that once we include other ‘not all’ implicatures, for *some*, for example, implicature generation may turn out to be more frequent than nongeneration. First, contra current aspirations for as elegant a grammar as we can devise, I am not committed to imposing uniformity of analysis on different expressions.

The reason is that wherever possible, we should opt for a ‘just that’ lexical meaning, i.e., one where “what you see as the actually conveyed meaning – ‘more than half, but less than all’ – is what you get – as the lexical meaning”. Last, as I argue below, it is not just that pragmatics cannot account for the upper bound, and that it should not, because of the reasons just enumerated. I claim that the upper bound on *most* should be lexically specified, because questionnaire results (initially reported in Ariel 2003, but see also 2004 and 2.2 below) attest that the upper bound on *most* is lexically given.

The structure of this paper is as follows. Section 2 presents my lexical analysis for *most*, and cites two different types of arguments for it (2.2, 2.3). I then address possible difficulties with the proposal, namely, *most*’s compatibility with ‘all’ (3), and the alleged difference in unilaterality between numbers and scalar quantifiers (4). Section 5 discusses the possibility of doing away with the lower bound on *most*. We conclude with section 6.

2 A ‘just that’ semantics for *most*

2.1 A new lexical meaning for *most*

I propose below that the set of entities that *most* picks for predicating on, as specified by its coded meaning, includes any value larger than half, but smaller than all (51–99% for short). There is, of course, no controversy about the lower bound of *most* (but see section 5 below). The question concerns the maximal value that *most* can code. I propose that the maximal value for the entities picked for predication is 100% minus something. While I am proposing a(n upper) bounded range analysis for *most*, unlike the bilateral analysis, I assume that the coded meaning of *most* does not specify anything about the complement minority. It certainly does not entail that the predicate does not apply to ‘all’.⁴ The complement may or may not then share the property predicated on the majority set, in which case it could have theoretically formed part of the same set in reality (as allowed for by the unilateral approach – see section 3 below).

When the speaker says *Most Israelis decided for peace*, *most Israelis* represents a set of Israelis, which constitutes 51–99% of another, larger set of all Israelis, and it is about this proper subset that ‘decided for peace’ is predicated on. The complement minority set of Israelis not included under *most* is not denoted at all, and at the lexical-semantic level, the speaker undertakes no commitment regarding the application of the predicate to these Israelis: If they are not referred to by the speaker, she is (lexically) committed neither to the predicate not applying to ‘all’ nor to the predicate possibly applying to ‘all’. This is where my analysis is crucially different from the standard semantic analysis of *most* (see Keenan 1996), which commits the speaker to relying on knowledge regarding whether or not the predicate (e.g., ‘decided for peace’) is or is not true

Equally controversially, I will not be surprised if we eventually find out that all scalar expressions are lexically circumscribed, just like *most* is.

⁴This is why the following, from Horn (2005:ex. 49a) does not argue against my analysis:

- (i) I’m surprised that **most** of the students failed the test.

Since the interpretation of (i) is not the speaker is surprised that ‘**not all** the students failed ...’, he concludes that ‘not all’ is an implicature rather than part of the semantic meaning. But of course, my analysis has an upper bound, not an ‘all-exclusion’ component in the lexical meaning, so I don’t predict what he attributes to me. Rather, on my analysis, the speaker is surprised that the profiled reference set (extensionally, between 51–99% of the students) failed.

for all Israelis. On the standard analysis, it's either the case that the speaker knows that there definitely are Israelis that did not decide for peace (but there are less of them than there are of Israelis who did decide for peace), or that there are none. First, I do not think that these assumptions are realistic. It's unrealistic to expect speakers to know what exactly counts as *decide for peace*, and who exactly it is true/false for. Second, and even more importantly, the account fails to distinguish between the profiled reference set and the nonprofiled complement, since both figure in the definition. Figure 1 may help the reader visualize what I have in mind:

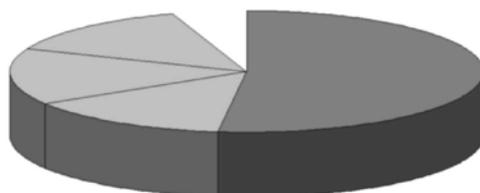


Figure 1: The lexical meaning of *most*

Note that more than half of the ellipse is highlighted. The bare minimum (lower bound) is colored dark gray, and the rest of the potential *most* range is colored light gray. The complement of *most* (the white space) is not only not highlighted, it's simply not there. Nothing is predicated on it, because it is not being denoted. At the **lexical** level, the status of the complement set is similar to that of any irrelevant referent (e.g., *some Americans* when (1) is uttered). The meaning I am proposing for *most* is indeed bilateral, but it is not identical to the bilateral interpretations assumed in the literature (neither the Hamilton 1860 style lexical meaning nor the commonplace conveyed meaning). 'All' is not included under the predication, but neither is 'not all'. This is why I prefer the term 'circumbounded' over 'bilateral'.

What I am proposing for *most* parallels Kadmon's (1987; 2001:68–76) analysis of the meaning of numbers for the most part. Kadmon argues that *four* codes 'a set of four', rather than 'at least four'. The former meaning is nonetheless compatible with an 'at least four' reading (i.e., where 'five' is true), because even though the NP stands for a set that contains exactly four entities, it is possible that there are additional entities bearing the same property outside the set.⁵ An argument adduced by Kadmon is that under an 'at least' analysis for the numbers, it is not clear why *at least four* would not also convey 'exactly four' (with the addition of a scalar implicature). After all, implicatures are generated based on semantic meaning, contextual assumptions and Gricean Maxims (due to nondetachability). Since under the unilateral analysis, the semantics of *four* and of *at least four* is the same (namely, 'at least four'), assuming a constant context and Grice's Maxims (invariant, of course), we should produce the same reading for the two expressions. But of course, this is not the case.⁶ Kadmon herself explains the difference by noting that adding on the scalar implicature in the *at least four* case would still not result in an 'exactly'

⁵I thank Nirit Kadmon for calling my attention to the similarity between my proposal for *most* and her analysis of the cardinal numbers. Kadmon, however, does not distinguish between the bounded lexical meaning of 'a set of exactly x' and the 'at least x' reading, because the latter is truth-compatible with the former. Still, unmodified by a scalar implicature, the coded meaning she proposes for the numbers gives rise to a bounded interpretation (just like *most* under my analysis), since it profiles a set of exactly n entities.

⁶The only way to distinguish between the implicatures of identical semantic meanings is by reference to the Maxim of Manner. *Four* and *at least four* can be seen as differing in brevity, the longer form reserved for the more marked

reading (for even if there are no additional entities bearing the same property outside the set, the set itself is not restricted to four). This argument applies to (*at least*) *most* as well. *Most* too must be distinguished from *at least most*, as in the following:

- (5) But, I believe that most authors are counting on **at least most** of the audience having some common points of reference (www.sondheim.com/commentary/collective.html, 1.18.2002).

Similarly, we can argue, following Koenig (1991) (who restricted his point to the numbers), that only by assuming a circumbounded meaning for *most* can we analyze *at least most* and *at most most* symmetrically and compositionally, even if they are rather rare in discourse (Koenig attributes this argument re the numbers to Kay 1989, and see also Geurts 1998).

How should we characterize this extension of *most*? One possibility is to say that *most* codes 'close to all'.⁷ Since 'all' and 'none' are obvious landmarks for assessing quantities, 'close to all' means that the quantity denoted is any majority, starting with the minimal 50% plus something.⁸ I tend to prefer an alternative definition, according to which *most* means 'a proper subset which is the largest subset, given any partitioning of the complement set (into one or more subsets)'. Like 'close to all', this meaning too embodies both the upper and the lower bound. The upper bound is secured since the set denoted by *most* is a proper subset (so *most* cannot denote 'all'). The lower bound is in place since the denoted set must be the largest subset, no matter how we partition the nondenoted complement. In order to constitute the largest (proper) subset, regardless of how many complement subsets there are, *most* must denote a set larger than 50%, for only such a value will be the largest even if the complement of the denoted set is taken as a single set (If the complement is taken to constitute two or more subsets, a lower percentage than 51% may constitute the largest subset, a possibility excluded under this analysis). For an alternative formalization of my proposal, see Appendix A.

This definition has a few advantages. First, it shows the current meaning to be a narrowing down of the historical meaning (something like 'the **greatest** part' Cf. French *la plupart*), and it also connects with superlative *most* ('to the **greatest** extent'). The semantic change is simply the addition of 'given the partitioning of the complement into **any** number of subsets'. Alternatively, the semantic change can be seen as an imposed reduction of any partitioning to a binary one. If there are only two subsets (one for which the predicate is asserted to hold, the other, its complement), then the larger subset must constitute at least 51%. The second advantage of this characterization is that with pragmatic loosening we can naturally account for certain uses of *most* as plurality, where it denotes the largest subset even when its value is lower than 51% (see section 5 below).

Assuming a circumbounded meaning for *most* makes for a more natural coded concept. Koenig (1991) finds it strange to assume that all languages choose not to code specific (circumbounded) numbers (as opposed to lower-bounded-only numbers). I find it equally strange to

reading ('at least four'). But then, the intuition is that it is the meaning of *at least* which contributes to the 'at least' meaning, and not merely the fact that it is longer.

⁷Of course, this does not mean that *most* and *close to all* are paraphrases and discursively interchangeable. As Bernard Comrie (p.c.) notes, *close to all* is restricted to large majorities.

⁸Note that some researchers consider *close to* or *near* to also include identical (in location). For example, Levinson (2000:96) assumes that *X is at the station* entails *X is near the station*. I do not think that that is the case in general, and do not intend *close to* or *near* to be interpreted in that way.

assume that all languages choose not to code circumbounded range quantifiers (such as ‘most’). ‘At least most’ is not as natural a concept to be lexicalized by a monomorphemic lexeme (Cf. the complex *more than half*, *at least half*). Thus, even if it does not make any wrong predictions, the burden is on those who wish to support the lexical unilateral meaning to show why it is necessary to assume such an abstract, counter-intuitive meaning. All the more so if, as I believe, the circumbounded assumption accounts for facts which the received view fails to account for (the discourse data analyzed in Ariel 2004, the questionnaire data discussed in 2.2 and 3.2, and the wise-guy examples in 2.3).

Next, the circumbounded meaning of *most* does not feel as a derived, indirect interpretation, split into a lexical component vs. a pragmatic enrichment component. Rather, it feels as one unified meaning. This intuition is shared by some unilateral analysis proponents. Geurts (1998:105) repeatedly states that “we simply do not know if scalar expressions license bilateral interpretations on a regular basis”, and that not many contexts make the bilateral interpretation relevant (he discusses *bright* in this connection). In other words, Geurts, who does not endorse my view, feels that the ‘all’-exclusion implicature is not often generated. I doubt he can then support the implication from his claim, namely, that in most cases it is the unilateral (‘at least’) meaning which is relevant (see the data in Ariel 2003, 2004). But if so, how can the interpreted upper bound come about on the received view? Chierchia & McConnell-Ginet (1990:194), again proponents of the unilateral view, also undermine the implicature analysis when they say that “scalar implications . . . often seem completely direct in the sense of being **part of what the speakers mean**” (emphasis added). I think that Geurts’ intuitions are right on target for the ‘all’-exclusion interpretation (it is indeed a rare pragmatic implicature). And I think that Chierchia and McConnell-Ginet’s intuitions are quite correct about the upper boundedness of *most* (this is the unmarked **conventional** interpretation of *most* – see Ariel 2004).

The main advantage of my analysis of *most* is that it is a ‘just that’ semantic analysis. This contrasts with the received view, which I find to be a ‘surplus semantics’ one, where a surplus interpretation (‘possibly all’) is included, only to be eliminated in the majority of cases. Under a ‘just that’ analysis, *most*’s prevalent interpretation (upper-bounded, but not necessarily ‘all’-excluding) is assumed to be its lexical meaning. Its compatibility with ‘all’ is not automatically guaranteed, and requires pragmatic mediation (see section 3). I should emphasize, however, that ‘all-exclusion’ and ‘possible all-inclusion’ are both potential Particularized Conversational Implicatures, generated in a minority of the cases (see Ariel 2004). I now present empirical evidence for my claim that *most* is lexically upper bounded (2.2, 2.3).

2.2 Questionnaire results

In order to help adjudicate between the received unilateral view and my circumbounded proposal, I designed a questionnaire whose goal was to determine whether *most* is lexically upper-bounded. To establish this, I compared people’s understanding of *most* with their understanding of the unilateral term *more than half*. Under the received view, both expressions are lexically unilateral, whereas on my view, *most* and *more than half* have distinct lexical entries. In particular, my view predicts that subjects should reject the possibility that by saying *most* the speaker could have intended ‘all’. The received view predicts that subjects should be reluctant to confirm ‘all’ because of the default scalar implicature, but that when pressed to cancel it, they will accept an ‘all’ interpretation. This is precisely my prediction for *more than half*, which pragmatically disfavors high values, ‘all’ in particular (see Ariel 2003 and Figure 2). Subjects should be reluctant

to select 'all' as *more than half*'s interpretation, but when pressed, will do it. On the received view, speakers using *more than half* do not generate the 'not all' implicature, and hence, they should have no problem accepting an 'all' interpretation for it.

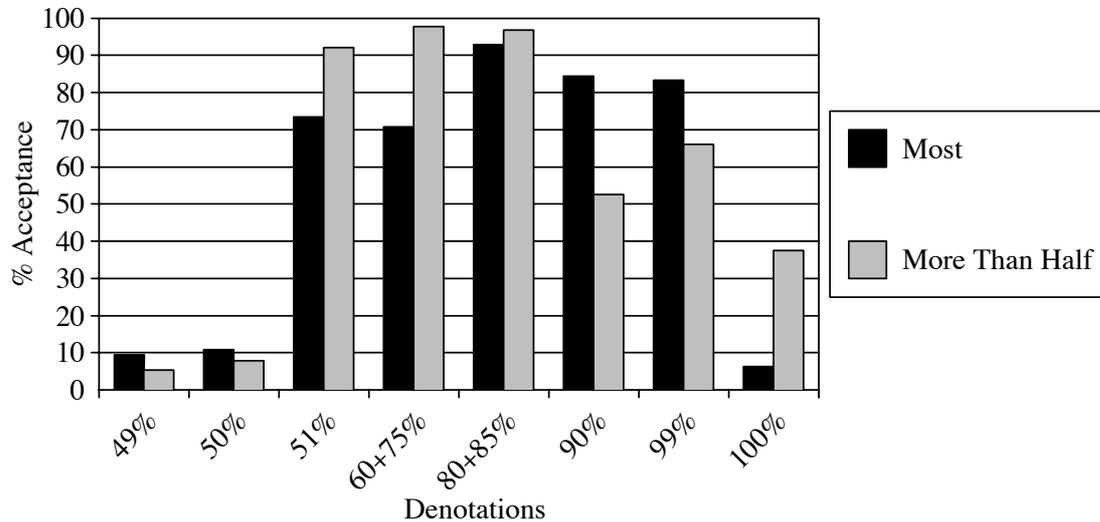


Figure 2: *Most* vs. *More Than Half*

I constructed a questionnaire in Hebrew, in which subjects were asked to determine whether certain values (e.g., 49%, 50%, 51%, 60%, 80%, 90%, 100%) were possible extensions of *most*. One group of 32 subjects was asked about Hebrew *rov* 'most', and another (19 subjects) about Hebrew *yoter mexaci* 'more than half' (for a detailed analysis of the results see Ariel 2003). (6) is one such question (each group receiving only one version):

- (6) **Most/more than half** of the students in the class were born in 1970.
How many students could the speaker mean?
- 100% of the students.
 - 20% of the students.
 - 50% of the students.
 - 49% of the students.
 - None of the above.

The results for the 3 questions including a 100% value are very clear (see Figure 2). The difference between the acceptance rates for 99% and for 100% as values for *most* is absolute, in fact, more so than for the lower bound, between 50% and 51% values. The great majority of subjects (93.75%) failed to confirm that the speaker of *most* could possibly intend a 100% value, just like they did for the 50% value (89.1%), even though many of them confirmed that the minimally different 99% and 51% are possible extensions of *most*. Moreover, the fact that subjects refrained from accepting 100% values for *most* cannot be explained by some absolute (pragmatic) aversion to the maximal value, because an impressive minority of the subjects did accept the 100% value for the Hebrew *more than half* (37.5%). Crucially, a pragmatic tendency cannot explain the different acceptance rates for the 100% value between *most* and *more than half*, for in general, (and

this is maintained up to the 99% level) it is *most* which tends to pick higher values (see Ariel 2003, 2004 and Figure 2).⁹

Obviously, questionnaire results can directly attest only to conveyed meanings, whereas I was interested in getting at the lexical meaning. I believe I overcame this obstacle, first, by not only instructing the subjects on the written form to circle **all** legitimate answers, emphasizing that even remotely possible ones were to be included, I also repeatedly emphasized it orally, even as they were filling out the questionnaire. Here's what the written instructions specified:

- (7) ... It's possible that **several** of the answers are appropriate. In such a case you should choose **all** the answers that the speaker might have considered possible, **even if chances for it are slim in your opinion** (original emphases).

Second, I always compared the results for *most* with those for *more than half*, which was used in exactly the same context. Note that an 'all' avoidance is attested for *more than half* as well. A majority of my subjects did resist an 'all' interpretation for *more than half* – 62.5% of them. Indeed, such a rejection rate represents a strong pragmatic tendency. A 93.75% rejection rate, however, reflects an absolute rejection.¹⁰ Third, I included questions where subjects were forced to suppress their pragmatic preferences. Question (6) cited above is one such case. It is then especially instructive to examine the results for it separately. If the received view is correct, then there should have been a substantial rate of acceptance of the 100% value for this question. First, based on world knowledge, it is quite plausible that all the students in some class were born in the same year. Second, options (b), (c), and (d) all violate the semantics of *most*, and it is quite clear that my subjects were quite reluctant to violate the semantics of *most*. According to the received view, all they had to do was cancel the 'not all' implicature.

Note that my method of forcing subjects to suppress their pragmatic practices was quite successful in general. For instance, when I posed a *most* question with the following values as answers: (a) 0%, (b) 15%, (c) 51%, (d) 46%, (e) none of the above, most subjects selected the 51% value (26/32, 81.3%), even though 51% is clearly a dispreferred pragmatic value for *most* (chosen by only 65.6% when a better option was also available, see Ariel 2003). This was so despite the fact that subjects were encouraged to circle as many options as they could. Similarly, when the potential answers for *more than half* and *most* were only (a) 97%, (b) 98%, (c) 99% and (d) 100%, subjects were willing to assign both expressions very high values (e.g., the rate of acceptance of 97%–99% was 88.9% for *more than half*), even though in other questions these high values received significantly lower rates of acceptance for both expressions (especially for *more than half*). Thus, the subjects I tested could certainly suppress their pragmatic preferences when forced to.

Given that subjects can be forced to select pragmatically dispreferred answers when no semantically viable answers are provided, one would expect subjects to choose (a) for (6), assuming that the speaker intended to cancel the 'not all' implicature, but maintaining the semantic lower boundary. After all, pragmatic implicatures are supposed to sometimes get cancelled, whereas

⁹Hans Smessaert (p.c.) also notes that *most* is "stronger", as he puts it, than *more than half*, so that one can say \sim *more than half* if not *most*, but not \sim *most* if not *more than half*.

¹⁰One of my referees objects to my saying that these rates constitute an absolute rejection of the 100% value, because it's not rejected in 100% of the responses. However, if we require 100% responses, I'm afraid we'll never find anything solidly semantic (see Table 1 below, where over 10% of the responses violated the semantic lower bound on both *most* and *more than half*).

semantic meanings are assumed not to. Indeed, for the counterpart *more than half* question in (6), many subjects did just that. 8/19 (42.1%) chose the 100% value, even though *more than half* pragmatically disfavors high majorities in general, and 100% in particular. This is not at all what I found for *most*. First, a majority of 26/32 (81.3%) chose answer (e) (Cf. 57.9% for *more than half*), equally rejecting the semantic violations (20%, 49% and 50% values) and the supposedly semantically legitimate option (the 100% value). This is a very high proportion of subjects choosing ‘none of the above’. For a similar question where the options were: 30%, 99%, 40%, 0% and none of the above, only 7/32 (21.9%) chose ‘none of the above’, even though 99% too is a pragmatically dispreferred value. The same happened for a question where the pragmatically dispreferred 51% was the only semantically viable option. Only 6/32 (18.8%) opted for ‘none of the above’ option here.

Now, although 3/32 subjects (9.4%) did select a 100% value in (6) for *most*, the same number of subjects actually accepted 49% (2/32, 6.25% also accepted 20% as a value). Even the semantically illicit 50% responses were more frequent than the 100% responses (5/32, 15.6%). All in all, 10/96 (10.4%) answers selected some semantically impossible option (20%, 49%, 50% values), as compared with 9.4% selecting 100% values. In other words, there were about the same number of semantically inappropriate responses as “pragmatically” inappropriate responses. For *more than half*, on the other hand, while a similar ratio of responses selected were semantically inappropriate (10.5%), a much higher ratio of responses were pragmatically inappropriate, (42.1%), 300% more. It seems therefore that for *most* (but not for *more than half*), the acceptance rates for the 100% value, even in the most favorable conditions, are as low as the acceptance rates for the ungrammatical lower values (below 51%). Table 1 gives the relevant data for question (6):

Violation type	<i>Most</i>	<i>More than half</i>
Semantic, lower bound (20%, 49%, 50%)	10/96=10.4%	4/38=10.5%
“Pragmatic” (100%)	3/32=9.4%	8/19=42.1%

Table 1: Rate of semantic and pragmatic violations for *most* and *more than half*

Clearly, one and only one result stands out in Table 1: The acceptance rate for a 100% value for *more than half*. The other acceptance rates are virtually the same. It is unreasonable that the same pragmatic violation is rather frequent in one case (the acceptance of 100% with *more than half*), but quite marginal in another case (with *most*). It’s even more amazing that whereas the so-called pragmatic violation for *most* is so different in frequency from the pragmatic violation of *more than half*, it is nonetheless identical in frequency to the semantic violations of both *most* and *more than half*. A more plausible analysis would group the three very low results together as reflecting the same, semantic ban, in contrast with the fourth, pragmatic result which stands out. Under my analysis, all semantic violations are marginal (9.4–10.5%), whereas the one pragmatic violation is not (42.1%). The conclusion must be that the 100% value has the same status as 49% and 50% values for *most*: It is ruled out lexically.

Laurence Horn (2005) counters that only *most* stands in pragmatic opposition with *all*, hence the strong avoidance of 100% values for *most*, but not for *more than half*. I doubt this. *More than half* too stands in contrast with *all*, as can be seen from the difference between the rates of acceptance for the minimally different values 99% and 100% for *more than half*. While the drop from 99% to 100% is by far smaller for *more than half* than for *most*, it is still a large difference (66.1% vs. 37.5% acceptance rates – these are the overall results for 100% acceptance

– see Figure 2). In other words, the majority rejected ‘all’ as an interpretation for *more than half*, and they did so far more often than they rejected 99% as a value. Even when 100% was the only semantically viable answer (as mentioned above), 57.9% of the subjects refrained from selecting it as a potential value for *more than half*. These results demonstrate that *more than half* too is pragmatically taken to stand in contrast with *all*. Yet, this pragmatic preference is overcome in a substantial minority of the cases. The counterpart ban for *most*, however, is not similarly lifted. This differential behavior is unaccounted for on the received view.

The results from the following, different set of questions, testing the truth-compatibility of *most* and *more than half* with ‘all’, testify to the same phenomenon (see 3.1 below for the difference between coded meaning and truth-compatibility):

- (8) Galit’s neighbor is giving a big party tonight, and Galit told him that she is willing to lend him **most/more than half** of the chairs in her house. “They’re in the dining room. You can come (pick them up) in the afternoon already”.

The neighbor knows that Galit is willing to lend him **all** the chairs in her house. How is he to interpret her words “They are in the dining room”?

- A. **All** the chairs that Galit has in her house are in the dining room.
 B. **Most/more than half (but not all)** the chairs that Galit has in her house are in the dining room.

This is a context where implicature cancellation is expected, since we are told that the reality is that Galit is willing to lend **all** her chairs. According to the received view, we then expect confirmations of ‘all’. Still, only 7/12 (31.8%) of the subjects asked about *most*, chose Answer A, and similarly, 4/11 (36.4%) of the subjects asked about *more than half*. Crucially, the majority chose Answer B for **both** *most* and *more than half*. This means that a similar majority avoided the 100% value for both *most* and *more than half*, even when they were told that ‘all’ was the case. Even in the one question where there was a difference between the two expressions, and *more than half* was seen as compatible with ‘all’ 1.75 times more than *most*, only 63.6% of the subjects confirmed that *more than half* is compatible with ‘all’. Clearly, *more than half* too pragmatically disfavors ‘all’.

Finally, consider another implicature generation + cancellation case (this question is part of a later questionnaire in Hebrew I administered):

- (9) Iddo: Dana solved **all** the problems.

Maya: **More than half** of them.

This is a case where the ‘not all’ implicature does not depend on *more than half* participating in a Horn scale. Rather, due to the opposition between *all* and *more than half*, Maya is implicating ‘not all’. Is Maya’s claim true/did Maya tell the truth in case Dana solved all the problems? Since ‘not all’ is only implicated, we would expect subjects to say that Maya’s proposition is true. Indeed, two thirds (14/21) said that ‘Maya’s claim was true’ and almost half (7/15) even chose the same answer for the question, ‘Did Maya tell the truth?’. On average for the two versions, only 27.8% (10/36) said that it wasn’t true (Others chose ‘impossible to know’). Now, if *most* too only implicates ‘not all’, we should expect similar results for the counterpart *most* utterance by Maya. But in fact, results are reversed for *most*. 30/38 subjects (79%) who responded to either version of

the truth question determined that Maya's contribution was not true (only 4, 10.5% said Maya's claim was true, and another 4 chose 'impossible to know'). All in all, subjects chose 'not true' for *most* 2.8 times more than for *more than half*, and they chose 'true' for *more than half* 5.5 times more than for *most*. Now, the two cases (for *most* and for *more than half*) are pragmatically similar, both triggering an 'all' exclusion interpretation: In both cases Maya is seen as offering a counter-proposal, replacing Iddo's universal claim with a partial generalization. Note that in this case Horn cannot claim that unlike *most*, *more than half* simply does not participate in a Horn scale, and hence, does not trigger a 'not all' implicature, because the context makes it clear that the speaker is generating the scalar implicature from *more than half* here. The different effect of the 'not all' implicature must therefore be due to the different lexical meanings of *most* and *more than half*. I suggest that whereas in the case of *more than half* there is some incongruity between the lexical meaning (not upper bounded) and the pragmatically appropriate interpretation ('all' excluding), for *most*, there is no conflict. The lexical upper bound is simply strengthened into an 'all' exclusion commitment. In other words, I am proposing that the differential responses derive from the difference in the lexical meanings of the two expressions, and not from a different pragmatics. Since for *more than half* 'not all' is only pragmatically derived, it is relatively easy to cancel, but since for *most* 'not all' is compatible with its upper bounded lexical meaning, it is not easily cancelable. These findings attest once again that the upper bound is different for the two expressions. There is a rather strong pragmatic tendency to avoid seeing *more than half* as denoting with 'all', but it is only a pragmatic tendency.

Another possibility to explain the difference between *most* and *more than half* regarding the 100% value is to invoke Horn's Division of Pragmatic Labor, and reason that since the speaker used the lengthier/marked form for the same semantic meaning, the lengthier form is used for the values for which the more economical *most* tends not to be used (Papafragou & Schwarz to appear). Indeed, when we compare the rates of acceptance for the 100% value, we do see a large difference between *most* and *more than half*: There is a 500% difference between *more than half* and *most* here. Now, if this is right, there is no reason to believe that the length difference is only relevant for one value, and we should find such a difference consistently for all values, for it is always the case that *more than half* is lengthier than *most*. This is not the case, however. I have found differences in the rates of acceptance for very low and very high values for the two expressions, where *most* favors the higher values and is dispreferred for the lower values (see Figure 2). But first, these are not at all as dramatic as the difference between *most* and *more than half* concerning the 100% value. Second, there is no difference between the two expressions for 80% and 85%. If *most* is favored for these values, and it is (92.9% of the responses accepted these values for *most*), how come so is the lengthier *more than half* (acceptance rate 96.8%)? Why isn't the length difference relevant here?

Finally, Horn (2005) attempts to explain the difference in the acceptance of 51% and 100% values between *most* and *more than half* by reference to what the relevant issue is. With low majorities, he reasons, the question is whether a majority obtains or not. *More than half* is then more appropriate, because it makes the 'above half point' salient. We agree on that. For the 100% value, on the other hand, he assumes, it is the totality which is relevant, and hence *most* is barred from denoting it, because it (and only it) contrasts with *all*. First, as I have argued at length in Ariel (2004), it's a rare case that the totality is relevant in discourse when *most* is used. Second, as I have just argued, *more than half* too forms an opposition with *all*. Third, Horn's assumptions here predict that subjects' choice of 51% and 100% values should be similar for *most*. Both values are semantically viable, even though they are pragmatically dispreferred.

This prediction, however, is not borne out. As mentioned above, in the question where I tried to force subjects to accept a 51% value, a substantial majority went along and accepted it for *most* (81.3%). This rate is 8.7 times higher than the rate of acceptance for the 100% under similar conditions (9.4%). In other words, a pragmatic preference is violated 770% more in one case than in another. I would say instead that the high violation rate points to a pragmatic violation, whereas the marginal rate of violation points to a semantic violation.

In fact, it does not seem to me that subjects were considering a comparison between *most* and *more than half* in my questionnaire, as is implied by Horn's suggestion (recall that each subject saw only one of these expressions). On my account, the different acceptance rates between *most* and *more than half* are pragmatic for the low (51–75%) and for the very high majorities (90, 99%), because whereas *more than half* makes the half line an explicit reference point, *most* tends to be used for a noteworthy quantity (see again note 1 and Ariel 2004). The difference concerning 100%, on the other hand, is lexical (only *most* is lexically upper-bounded). This is why it is so much stronger (14.1 times larger). There is, however, nothing to push for differences concerning the high (but not very high) values (80, 85%), and hence, no difference is found. In other words, contra Horn's prediction, there is nothing close to a cross-the-board division of labor between *most* and *more than half*, based on their formal or content differences, except for when the 100% value is concerned (see Figure 2, and Ariel 2003, 2004 for actual percentages).

My conclusion is that subjects absolutely refuse to accept that speakers could have meant 100% values when *most* is used. Since no pragmatic explanation can account for this absolute finding, I conclude that it stems from the lexical meaning of *most*, which is (lower- and upper-) bounded. Given the findings in 2.2 and in 2.3 below, it seems that not only is it the case that pragmatics cannot be responsible for the upper bound placed on *most* (the conclusion I drew in Ariel 2004), semantics is actually the appropriate linguistic competence to account for it.

2.3 An argument from 'wise-guy' interpretations

In Ariel (2002), I argued against an 'at least' coded meaning for the numbers, based on what I term 'wise-guy' interpretations. 'Wise-guy' interpretations are interpretations an interlocutor can insist on despite the fact that they are contextually inappropriate. The argument is that the 'wise-guy' can insist on a contextually inappropriate meaning if it is lexically coded, but not if it is only pragmatically derived. In other words, coded meanings can, but pragmatic meanings cannot overrule contextual factors. The store manager in (10) is such a 'wise-guy':

- (10) "A young couple went into the Allegro record store and offered to sell two CD's. The store manager offered the couple 40 sheqels. The guy, who looked like a Kibbutznik, said that in the store across the street he can get 50 sheqels. The manager of the store said that not on his life will he get such a sum. They took a bet ... The guy ... sold the CDs and got 55 sheqels for them. He took a receipt and went back to Allegro. Sorry, said the manager, you lost. I said you won't get 50 sheqels, and indeed, you did not get such a sum. I got more, explained the astonished Kibbutznik, but the sales woman laughed in his face" (A story in the magazine *Hair* 3.9.1990, originally Hebrew).

The store manager insists on a contextually inappropriate interpretation whereby 50 only denotes '50', when an 'at least' reading is contextually appropriate. The only reason she can get away with it, I argued, is that the circumbounded reading of the numbers is their coded meaning. The

'at least' interpretation is a pragmatic enrichment of that meaning. I contrasted (10) with a 'wise-guy' interpretation of a contextually inappropriate 'at least' reading, which is not successful. While the 'at least' interpretation can certainly be a legitimate contextual interpretation in some contexts (as in (10)), it is not entrenched (nor lexicalized) so as to allow interlocutors to insist on it in a context requiring a circumbounded reading:¹¹

- (11) ~Income tax clerk: How much money did you make last year?
 Taxpayer: \$10,000.
 Income tax clerk: Our information shows that you made \$15,897!
 Taxpayer: ??That's what I meant. I meant that I made at least \$10,000, and possibly more.

The same argument applies to *most* ('at least most' vs. 'circumbounded most'). When we substitute *most* for *50* in (10) above, we get the same 'wise-guy' effect of an extremely uncooperative store manager, who can nonetheless get away with a contextually irrelevant interpretation, just because it is an unenriched conveyed meaning (based only on the circumbounded lexical meaning). The 'at least' meaning is only derived:

- (12) ~"A young couple went into the Allegro record store and offered to sell four CD's because they needed 100 sheqels to repair their CD player. The store manager offered the couple 40 sheqels. The guy, who looked like a Kibbutznik, said that in the store across the street he can get **most** of the repair money. The store manager said that not on his life will he get such a sum. They took a bet . . . The guy . . . sold the CDs and got 100 sheqels for them. He took a receipt and went back to Allegro. Sorry, said the manager, you lost. I said you won't get most of the repair money, and indeed, you did not get such a sum. I got more, explained the astonished Kibbutznik, but the sales woman laughed in his face".

Since the store manager can insist on the circumbounded reading of *most* even though context calls for an 'at least most' reading, the 'wise-guy' interpretation principle suggests that it must constitute its coded meaning. Under my account, the store manager refuses to enrich the circumbounded meaning with 'at least'. Note that again, insisting on an 'at least' meaning of *most* in a context requiring the circumbounded interpretation does not create a 'wise-guy' interpretation, but rather, plain deception:

- (13) ~Income tax clerk: In how many of the past ten years did you fail to file your tax return?
 Tax payer: **Most** years.
 Income tax clerk: Our information shows that you failed to file in all those years.
 Tax payer: ??That's what I meant. At least most, and possibly all the years!

Since only the circumbounded meaning of *most* (but not the unilateral meaning) is a potential wise-guy interpretation, it must be its coded meaning.¹²

¹¹Following Chafe (1994:xiii), I mark constructed examples with ~.

¹²This assertion is based on the judgments of all 22 'Introduction to pragmatics (2002)' students. One of my referees does accept a wise-guy interpretation here, however.

A note is here in order. I have argued that only lexically coded meanings can be imposed when contextually inappropriate, but actually, the meaning imposed by the store manager in (10) & (12) seems to be ‘exactly fifty/most’. What is the status of ‘exactly n’? I maintain that the coded circumbounded meaning, when unmodified by pragmatic implicatures (scalar or others) is equivalent to the ‘exactly’ meaning for practical purposes (except that the ‘exactly’ aspect is still cancelable). While the circumbounded meaning is compatible with a higher scalar value (a higher number, or ‘all’ for *most*), it does not encode it, and addressees would have to consider additional entities outside the profiled set to get the ‘at least’ reading. Since they don’t normally do that, they are left with an ‘exactly n’ (or ‘exactly most’) interpretation.

We next discuss potential difficulties for the circumbounded analysis of *most*, truth-compatibility with ‘all’ (section 3), and the differential behavior of *most* and the numbers (section 4).

3 The compatibility of *most* with ‘all’

3.1 Truth-compatibility vs. lexical meaning

The main motivation behind the received semantic analysis of *most* as only lower bounded is that it straightforwardly accounts for the judgment that *most* is compatible with states of affairs in which ‘all’ is true. Even though ‘possibly all’ is the case in only 11/127 (8.7%) of the uses of *most* in my data see (see Ariel 2004), these examples are impeccable, and should therefore be accounted for.¹³ My subjects too sometimes assigned *most* an ‘at least most, possibly all’ interpretation when the context was thus biased (see (19) and section 4), and some of them also saw *most* as compatible with ‘all’ ((8), (18)).

However, creating a gap between a lexical meaning and the meaning commonly conveyed by it, via a ‘surplus’ semantic analysis, is clearly an undesirable step theoretically. Even if this gap can cleverly be bridged over by a scalar implicature, a ‘just that’ semantics is to be preferred, other things being equal. Hence, if we can otherwise account for the intuition that *most* is (sometimes) compatible with ‘all’, the motivation behind the received view will be drastically reduced. This is what I am suggesting here. My account for *most*’s compatibility with ‘all’ is based on the assumption that it is not false (only misleading, in many circumstances) to select a subset (‘most’) for predication when a larger set (‘all’) is true. This is not at all special to *most*.

My line of argument here builds on Koenig (1991). As Koenig (1991:140) points out (regarding the numbers only), we should not confuse two logically independent questions: (1) the fact that “Scalar predicates are (at least sometimes) treated discursively as logically **compatible** with a higher value on the scale they evoke” (applied to *most*, this would mean that *most* is sometimes compatible with ‘all’), and (2) “the theoretical claim that the **lexical meaning** of scalar predicates specifies only a lower bound on this scale” (emphases added) (applied here, this would mean that *most* means ‘at least more than half, possibly all’). There is then a difference between expression X being compatible with a certain state of affairs, and expression X coding that state of affairs.¹⁴ *Most* is certainly compatible with ‘all’ being the case (on some occasions),

¹³As corroborating evidence, consider the fact that especially children have been found to respond ‘true’ to statements containing *some*, which they know to actually be true of ‘all’ (e.g., *Some elephants have trunks*) (See Smith 1980, Noveck 2001, Papafragou & Musolino 2003). Note that since most experimental research has been conducted on *some*, rather than on *most*, and since the received view assumes a similar bilateral analysis for both expressions, I here use research on *some* as well.

¹⁴Smith’s (1980:199) automatic conclusion re the meaning of *some* from verification tasks hinging on being **com-**

but these scalar effects do not necessitate the received lexical meaning of *most*. In fact, they are language-independent. If we accept the distinction between meaning and compatibility with states of affairs, we no longer have to stipulate as part of the coded meaning 'surplus' aspects of reality which are compatible with the use of the expression. *Most* can then commit the speaker to asserting something about 'most but less than all' relevant entities, even though its use is compatible with the predicate being true of 'all'.

Of course, all analyses distinguish between coded meanings which pertain to relevant states of affairs, and irrelevant states of affairs compatible with such meanings. The controversy here is really about what constitute lexically relevant states of affairs for *most*. The received view sees 'all' as pertinent to what's asserted about the reference set, whereas I suggest that discursual usage shows that it isn't. On my analysis, 'all' and the complement of *most* are outside the discussion **lexically**, but this characterization does not preclude situations where *most* is truthfully used when 'all' is true, just because 'most' forms part of 'all', and there may be additional members outside the profiled set (i.e., in the complement) for which the predicate happens to hold. Parts are often (but not always) true when wholes are (see below).

Let's examine Koenig's arguments more closely. Koenig's thesis about scalar predicates seems to be the flip side of the Hirschberg (1991) coin. While she argues that nonconventional scales are like conventional Horn scales in terms of their implicatures, he argues that conventional scalar expressions should be analyzed like nonconventional ones in terms of lexical meanings. Recall that Hirschberg discusses examples with scalar implicatures in the absence of a conventional Horn scale, from which she concludes that scalar effects also occur for items not forming conventional Horn scales. Koenig relies on precisely such nonconventional cases to argue against the semantic analysis proposed for conventional Horn-based scalar predicates. Just like the ad hoc scales do not force us to assume a unilateral lexical meaning for the ad hoc scalar items (this would be a 'surplus' semantic analysis), so we should not impose a unilateral lexical meaning on items which do form part of a conventional Horn scale. Consider (14a). As Hirschberg argues, such examples can generate scalar implicatures (what is negated is the implicature 'only Catholics'), despite the fact that there is no conventional scale such as (14b) (see the examples in (28b-d):

- (14) a. \sim CATHOLICS are not oppressed, ALL MINORITIES are oppressed (Koenig's example # 7).
 b. <Minorities, Catholics>

Koenig relies on the fact that no lexical consequences follow from such examples (i.e., we would not suggest that Catholics means 'Catholics, and possibly other minorities') to argue that no lexical consequences need follow from similar examples involving conventional Horn scalar predicates either. Koenig, of course, only applies this argument to the numbers, but I don't see that the argument is restricted to the numbers. It equally applies to *most*.

The numbers are not the only linguistic expressions where truth-compatibility may be distinct from the speaker's coded meaning. Sevi (1994) argues that while the overwhelming majority of Hebrew *xuc me* ... 'except for ...' uses are such that the relevant predicate does not apply to

patible with a certain reality (where children had to indicate whether a statement such as *some elephants have trunks* was true or false) is typical: "children interpreted *some* as **meaning** *some and possibly all*" (emphases added). I believe that Papafragou & Schwarz (to appear) also take tests about truth-compatibility as attesting to meaning (see below).

the entity modified by the adverbial, this is merely a pragmatic inference, rather than a semantic entailment. Hence, it's not lexically specified for *xuc me* ... A speaker using the Hebrew counterpart of 'all except y' asserts some proposition regarding the set minus that entity, regardless of whether the predicate is or is not true of it (the entity has to be exceptional in some sense, but not necessarily in that the predicate does not hold of it). This lexical meaning, then, does not preclude the possibility that in reality, the predicate does apply to the argument. Indeed, in (15) 'won here' also holds for the argument modified by the Hebrew 'except for' (indifference):

- (15) Who won here, **except for** indifference? (Originally Hebrew, Kol Israel radio interviewer, 10.29.2003).

Having used *xuc me* ..., the speaker indicates that it is not his goal to assert that 'indifference won', because he doesn't want the interviewee to discuss this "winner". As a topic, 'indifference' is excluded, because it is not included under the predication of 'won here', but as a fact 'indifference having won' is not ruled out.

In fact, this distinction between meaning and truth-compatibility is relevant for any linguistic expression. For example, it is not false to predicate something only of mother, when the reality is that the predicate actually holds of 'both parents, and the rest of the family'. Note the following (the three Israelis named were believed at the time to be alive, captives of the Hezbollah in Lebanon):¹⁵

- (16) Benny Avraham, Adi Avitan, Omar Su'ad **Mother** is waiting at home (Originally Hebrew bumper sticker 1999).

Just like we are no longer forced to stipulate that *five* codes 'at least five', just because it is not false to say *five* when 'eight' is the case (see section 4), so too the fact that we can use *most* when 'all' is true should not force us to assume that the coded meaning of *most* covers 'all'. My formulation treats the complement of 'most' in the same way it treats 'father' when *mother* is used. Both are external to the (explicit) discussion. Both are lexically irrelevant.

Once we are willing to distinguish between the compatibility of expression X with certain states of affairs and X's meaning (provided this compatibility is either irrelevant or extralinguistically accounted for), it is hard to see what compelling evidence there is to viewing the lexical meaning of all scalars as unilateral. A 'just that' semantics should be preferred over a 'surplus' semantics.

3.2 Truth-compatibility as (partly) pragmatically determined

Thus far, I have taken for granted the received intuition, that *most* is consistent with 'all' (see Horn 2005). While this intuition is naturally accounted for on my proposal (because 'most' denotes a part within a whole, and what's true of some part, may in fact be true of the whole, even if the speaker did not refer to all set members), it seems that conversationalists do not actually interpret *most* as compatible with 'all' automatically. Rather, it depends on the context and on the addressee. Noveck (2001) finds that subjects vary on whether they judge sentences such as *Some elephants are mammals* as true or false. Papafragou & Musolino (2003), on the

¹⁵It is not surprising to learn that the children tested by Smith (1980), on the one hand did confirm that *some* ... is true when actually 'all' is, but, on the other hand, they frequently added a reference to *all* in their responses.

other hand, find that adults overwhelmingly reject such statements. I suggest that this points to the pragmatic nature of the compatibility with 'all' of *some* (and equally, of *most*). Unlike semantic judgments, pragmatic judgments are expected to be variable (see also section 4).

Since I hypothesized a difference between the coded meaning of *most* and its truth-compatibility with 'all', a different type of questions was included in the questionnaire mentioned above, in order to investigate the nature of the compatibility of *most* with 'all'. (18) is a variation on Horn's (2005:18) (47a), quoted here as (17), about which he says that "it is clear that once the Dolphins won all their games ... I won the bet":

(17) I'll bet you that the Dolphins will win **most** of their games this year.

In (18), subjects were not asked to decide what the speaker's possible intended meaning was per se. Rather, they were encouraged to presuppose that somebody is entitled to some prize for having made the correct guess (the question presupposes a winner), so that all they needed to decide was who was closer to the correct value, whose answer was more compatible with 'all'. Nonetheless, the majority of the subjects resisted my encouragement, and they resisted it equally for 80% and for *most*. (18) does not only demonstrate the similarity between *most* and the numbers (80% in this case – see section 4), but what is more crucial to the point at hand, it shows that the compatibility of *most* with 'all', while not ruled out by the semantics of *most*, must be established via pragmatic reasoning. Contra the received view assumption, speakers are far from unanimous in confirming it:

(18) The catering company manager announced that if anyone would guess how many of the guests would prefer square plates, they would win a dinner set. Dana guessed that **most** of the guests would prefer square plates, Oren guessed that none of the guests would prefer square plates, and Iddo guessed that **80%** of the guests would prefer square plates.

Question: At the end of the event, it was found out that **all** the guests preferred square plates. Who is entitled to the promised prize?

Answers: A. Dana B. Oren C. Iddo D. Nobody

All answers but one made one and the same decision about Dana (*most*) and Iddo (80%). 8/24 (33.3%) selected both as winners, and 15/24 (62.5%) chose 'nobody'.¹⁶ It is worth noting that contra Horn's intuitions regarding (17), almost twice as many subjects thought that *most* (just like 80%) is incompatible with 'all'. I had 4 truth-compatibility questions for *most* on the questionnaire, and 'all' acceptance varied between 5.9% and 83.3% (three of these four questions are cited in (8), (18), (19)). This variability points to the pragmatic nature of the part-whole inference here.

The results in (18) resemble those of (8), but here's a case where 'all' acceptance was contextually supported, and hence, accepted by most of the subjects (Knesset is the Israeli parliament):

¹⁶One of my referees comments that it's possible that the identical responses for 80% and *most* may be independent of each other. For instance, that 80% was rejected because it is different from 100%, but that *most* was rejected because it's not an informative enough alternative for a guessing game. I cannot rule out this possibility, but it is highly suspicious that 23/24 of my subjects made exactly the same decision about *most* and about 80% (in whatever direction). In other words, why should it be that exactly the same set of people who thought that guessing 80% does not entitle Iddo to the prize also think that *most* is an illegitimate guess, whereas precisely the set of people who thought that 80% entitles Iddo to the prize also thought that *most* is an appropriate guess?

- (19) It is not necessary for all 120 Knesset members to be present for the Knesset to convene. The law requires that **most** Knesset members participate in the assembly discussion tomorrow. **They** will be asked to vote on a series of social laws.

Question: Who are the “they” who “will be asked to vote on a series of social laws”?

Answers: A. All the Knesset members who will participate in the assembly discussion tomorrow, that is, no less than 61 and up to 120 Knesset members, including 120 Knesset members.

Or:

B. All the Knesset members who will participate in the assembly discussion tomorrow, that is, no less than 61 and no more than 119 Knesset members.

The great majority of the respondents chose answer A (25/30, 83.3%), assigning *they* an ‘at least most’ reading. This is quite different from the previous rates of truth-compatibility with ‘all’ (31.8%, 33.3%).

Varying the context also dramatically influenced the 5 year-olds tested by Papafragou & Musolino (2003). Whereas in their first experiment the children overwhelmingly accepted *some* when ‘all’ was the case, in a second experiment they tended to reject *some* when ‘all’ was the case (this is the adult pattern). The authors explain the different responses by reference to the changed context they created in the second experiment. The second context, they argue, more readily invites scalar inferences. The reason is that in the second experiment only ‘all’ was relevant. I propose instead that what the changed context did was block part-whole inferences more readily. If only ‘all’ is relevant, then stating *some* when ‘all’ is true relies on a valid whole-to-part inference, but is pointless. In general, then, what are taken as circumstances favoring the generation of a scalar implicature, I mostly view as circumstances blocking inferences about the compatibility of parts with wholes. What is viewed as circumstances blocking scalar implicatures I view as circumstances favoring/not blocking inferences from wholes to parts (or upward compatibility). Thus interpreted, the recent interesting findings re the appropriateness of *some* when ‘all’ is true cannot be used to argue against the circumscribed view.

Much the same point applies to Papafragou and Schwarz’ (to appear) attempt to justify a Neo-Gricean account for *most* over the analysis here proposed. In apparent contradiction to my findings (see 2.2 and see Ariel 2003, 2004), Papafragou and Schwarz find that many of their adult subjects accepted 100% as compatible with *most* – 56.7%. Although this acceptance rate is no different from chance in their experiment, it seems to be dramatically higher than the rate of acceptance for 100% among my subjects (6.25%).¹⁷ It is unlikely that the difference in languages (Greek versus Hebrew) accounts for this difference. Rather, I believe that what Papafragou and Schwarz tested for is the **truth compatibility** of *most* with ‘all’. Their questions are then comparable to the second set of questions on my questionnaire (exemplified in (8), (18) and (19) above), where the focus was on *most*’s compatibility with ‘all’, not on its possible extensions.

The analysis here proposed accepts that *most* may (but need not) be viewed as compatible with ‘all’ (see the relatively high rate of ‘all’ acceptance – 83.3% for the question quoted in (19)).

¹⁷Whereas on my questionnaire subjects could always choose a ‘none of the above’ answer, Papafragou and Schwarz’ subjects could only choose between confirmations and disconfirmations. Hence, the difference is actually not quite as large as it looks.

As I have tried to emphasize, however, truth-compatibility with 'all' does not amount to **coding** 'all'. This is why I posed two different sets of questions to my subjects: some on its possible extensions and others on its truth compatibility with 'all'. And indeed, as we have seen, the results are remarkably different. In the first set of questions I aimed at getting at the meaning of *most* by asking about all the (**even remotely**) **possible values** that the speaker could have in mind while uttering a *most* utterance. In the second set of questions I asked whether a *most* utterance can be taken as compatible with a **reality** where 'all' is the case. Papafragou and Schwarz, on the other hand, asked subjects about the truth-compatibility of *most* with 'all' (they asked whether the character "did ok" when they performed 'all', having said that they would perform *most*). Under such circumstances, I claim, part-whole inferences are less blocked, and subjects reasoned that performing 'all' counts as performing *most*. A later questionnaire I administered more clearly underscores the distinction between the two types of questions. Among subjects asked about the meaning of the Hebrew counterpart of *The teacher already knows most of the students*, only the usual marginal percent (2/23, 8.7%) chose a unilateral meaning (defined as 51–100%). However, a majority (11/15, 73.3%) of the subjects asked about the truth of the proposition in case the teacher already knows all of the students confirmed *most*'s truth-compatibility with 'all' here. The gap between the responses on this very same proposition is huge: 8.4 times more truth-compatible with 'all' confirmations than unilateral meaning confirmations.

An additional piece of evidence for my claim that 'all' confirmations for *most* are mediated by contextual assumptions comes from examining the breakdown of the responses to 100% in the Papafragou and Schwarz experiment. It turns out that 50% of their subjects (5) virtually always confirmed it (14/15 trials), but the other 50% (5) always rejected it (15/15 trials – see their note 10). Noveck & Posada (2003) similarly found (in two experiments) that 5 subjects confirmed sentences such as *some elephants have trunks* virtually always, whereas 7 subjects consistently determined that they were false. Such individual consistent differences point to the role of pragmatics in the process.¹⁸ All in all, the variability among subjects and across pragmatic contexts attests to the pragmatic rather than semantic basis for the truth compatibility of *most* with 'all'. Moreover, note that whereas in questions of truth compatibility both I and others have found a wide variation among subjects (e.g., Cf. Papafragou and Schwarz's truth-compatibility findings vs. mine, and the different results I got for different questions), I did not find such a variation for the questions on the possible intended meaning of *most*. On my original questionnaire, 29/32 subjects absolutely refused to confirm 'all' when *most* was used, even though sometimes, no other legitimate alternative was available.¹⁹ The uniformity in responses to the meaning type of question vs. the variability in responses to the truth-compatibility type of question supports my claim that the questions test two different sorts of interpretations.

Now, it may seem that questionnaire data simply cannot decide between the received unilateral view and my circumbounded meaning proposal: When subjects refuse to accept an 'all' interpretation I account for it by reference to subjects' refusal to allow part-whole inferences, whereas the unilateral theories account for it by reference to the 'not all' default implicature, and when subjects are willing to accept an 'all' compatibility, I explain it as subjects allow-

¹⁸The alternative of assuming lexical idiolects seems less attractive.

¹⁹Only one subject was consistent in accepting the 100% value in the three possible meaning questions. One accepted it in 2/3 questions and another accepted it in 1/3 questions. Interestingly, however, the subject who accepted the 100% value for *most* in 3/3 questions also accepted 20%, 49% and 50% values, and for the most part did not confirm that *most* was truth-compatible with 'all' (in 3/4 questions). In fact, she only confirmed the compatibility of *most* with 'all' for the question with a strong bias towards an 'at least' reading of *most* (19).

ing for part-whole inferences, whereas the unilateral view explains it by reference to the lexical meaning with implicature cancellation. Is there then no way to decide between the competing proposals? I think there is, because I don't think that the received view account is sound. It actually comes with a very high cost to the very Gricean theory it is couched in, specifically, to the concept of conversational implicatures. Recall that implicatures are defined as (i) cancelable and (ii) nontruth-conditional.²⁰ In their attempts to account for the empirical data, received view proponents have compromised both of these characteristics. In order to account for the first set of questions, on possible meanings, they in effect assume that implicatures are not cancelable. When truth compatibility with 'all' is not confirmed they don't only assume that implicatures are not cancelable, but also, that they determine truth conditions.

First, regarding the first set of questions, probing all possible extensions of *most*, I don't see how the unilateral view can account for why subjects absolutely refused to accept an 'all' interpretation, especially when no other interpretation was available. Received view theories assume that implicatures are cancelled under circumstances which render them false. Why didn't the subjects simply cancel the scalar implicature? Recall that on my "meaning" questions, subjects were constantly encouraged to (also) select highly unlikely answers (see the instructions in (7) again). This should have prompted the subjects to cancel the scalar implicature, which would have resulted in comparable confirmation rates for 100% values for these questions as for the 'all'-compatibility questions. Why does Horn assume that the implicature is cancelled in the betting context (17), but not in these questions? In fact, the 'all'-compatibility "cancellation" cases were never close to absolute, even in cases where the scalar implicature has been contextually cancelled (e.g., (8), (9), (18), and the numerous "elephant" experiments). The experiments mentioned above even more clearly involve such circumstances. Surely, subjects must have been aware that **all** elephants are mammals/have trunks. They should have then reasoned that the addressor couldn't possibly intend the scalar implicature, and the statements should have therefore been considered true, at least by a majority of the subjects. As mentioned above, very often they weren't, however. In order to explain such judgments, received view proponents have to assume that implicatures are not cancelled, even under favorable circumstances,

Second, it's not at all clear how theorists who assume an implicated upper bound can account for the rejection of infelicitous *some* (or *most*) statements as false, even if the implicature is not cancelled. After all, implicatures are supposed not to determine truth conditions. Indeed, when we test the effect on truth conditions of an interpretation which is uncontroversially an implicature, we see that its falsity does **not** render the implicating assertion false. Thus, in a separate questionnaire, one set of subjects was asked whether the Hebrew counterpart of the following suggests that proportionately more women or more men wore evening clothes:

(20) **Most** of the women and **more than half** of the men wore evening clothes.

As expected, many subjects (47.1%) confirmed the implicature that *most* denotes a more noteworthy majority than *more than half* (only 2% thought that proportionately more men wore evening clothes). However, when a different set of subjects was asked whether (20) was true in case 70% of the women and 85% of the men wore evening clothes, only one of them (4.5%) said that it was false. The majority (68.2%) determined that the statement was true. In other words,

²⁰Whereas Levinson (2000) is open to the possibility that implicatures affect the truth conditions of the proposition used to implicate them, Horn (2004) reiterates his commitment to the original Gricean position, except for special cases of reinterpretation.

false implicatures do not usually affect truth conditions. Why should the scalar implicatures be different then? Recall the differential truth evaluation for *more than half* and *most* in (9), where the scalar implicature for *more than half* did not affect truth conditions. It seems that on the received view scalar implicatures must be truth-conditional sometimes, for they do determine truth conditions (recall also that very often subjects find that *some* and *most* statements are false even though they know that ‘all’ is the case). Such assumptions cast doubt on the implicated status of the upper bound. If ‘not all’ is not cancelable, and if it determines truth conditions, it suspiciously resembles lexical meaning.

Note, in addition, that while received view proponents are quick to rely on the ‘not all’ pragmatic implicature to account for subjects’ truth value judgments, they are equally quick to ignore the scalar implicature when they consider other semantic phenomena, such as the effect of negation (see 4.1) and the assumed upward monotonicity of *most*. For example, why isn’t the negation of *most* taken as applying to its conveyed upper-bounded meaning, possibly denying ‘all’? If the implicature is allowed to influence truth value judgments, then the received view too (and not just the circumscribed view – see 4.3) has to account for why the scalar implicatures seem **not** to affect other semantic processes. They can’t have the cake of scalar implicature (invoke it for truth value judgments) and eat it too (eliminate it for negation). In conclusion, I am proposing that *most* only profiles a majority reference set as its lexical meaning. The predicate is asserted to be true of that reference set, and it is silent about the complement set. Contra the received view, its compatibility with states of affairs in which the predicate is true for ‘all’ is not lexically coded. It is in principle not ruled out semantically, but it must be extralinguistically established (via our context-sensitive assumptions about part-whole relationships).

4 *Most* and the numbers: A potential problem for the circumscribed analysis

Examples (10) and (12) show a similar interpretative pattern for the numbers and *most*, where the wise guy can insist on an upper bounded meaning, despite the fact that context calls for a lower-bounded-only interpretation. Similarly, the virtually identical interpretations of *most* and 80% in (18) suggest that we should perhaps not distinguish between *most* and the numbers. Now, according to Geurts (1998), the majority view for numbers is nowadays bilateral (but see Levinson 2000). Carston (1990, 1998) has argued for a ‘just that’ semantic analysis for the numbers, according to which *eight* means ‘eight’ (and not ‘at least eight’, see also Sadock 1984, Kadmon 1987, Récanati 1989, Koenig 1991, Geurts 1998, Ariel 2002). The appropriate conveyed meaning actually adopted (‘at most eight’, ‘at least eight’, ‘exactly eight’ or, I should add, ‘about eight’) results from a general process of adaptation of the coded meaning (via the Principle of Relevance) to create the ad hoc concept relevant for the specific context (see also Ariel 2002).²¹ Horn (1992, 1996, 2003) too now believes that the pragmatic enrichment creating the bilateral meaning (‘exactly eight’) forms part of ‘what is said’, i.e., having a (referential, truth-conditional) semantic status, unlike regular Generalized Conversational Implicatures. In other words, despite the fact that lower numbers are often compatible with states of affairs in which higher numbers are true, current analyses for the numbers are upper-bounded for the most part. Shouldn’t the same analysis then be applied to the scalar quantifiers as well?

²¹Carston (1998), however, is undecided between a semantic meaning of ‘exactly x’ and ‘x’ for the numbers, but see again my discussion at the end of 2.3 as a possible explanation for the centrality of the ‘exactly’ reading, despite the general lexical meaning.

Despite the shift in their analysis of the cardinal numbers, Sadock (1984), Horn (1992, 1996, 2003), Geurts (1998), and Carston (p.c.) explicitly choose not to apply their bilateral semantic analysis of the numbers to other scalar predicates. The reason is that various differences have been pointed out between the numbers and scalar quantifiers. If I am correct, then, and *most* does have a circumbounded rather than unilateral lexical meaning, it would seem that it should pattern with the numbers. Indeed, Horn (2005) uses differences between *most* and the numbers to argue against my proposal. In order to justify my analysis I must therefore either show that *most* does behave just like the numbers, or else, that the differences between them are orthogonal to the question of unilaterality vs. circumboundedness. These are the goals of section 4.

I will propose that the differences between *most* and the numbers are, first, far from absolute, which they should have been if they were semantic. Second, they are orthogonal to the circumboundedness question. I will argue that the differential patterns noted in the literature are derivative from a number of differences between the numbers and *most*, which are quite irrelevant to the question of lexical upper-boundedness. In fact, as we shall see below, once we neutralize these differences, *most* and the numbers do pattern similarly. In other words, *most* may manifest a “number, bilateral pattern” (as we already saw in 2.3 and in 3), and the numbers may manifest a “*most*, unilateral pattern”. If that is the case, then the fact that the numbers and *most* differ in some contexts is no hindrance to viewing *most* too as upper bounded.

I would actually like to try and take my conclusions even further than that. I would like to propose that the set of behaviors assumed to define ‘the unilateral pattern’, e.g., participation in downward entailments, upward compatibility, which allows for the smooth shifting to higher values and for ‘at least’ readings, nonco-occurrence with certain quantifiers, etc. consists of a number of independent propensities, each in response to some characteristic. While these properties often cluster together, so that some expressions in some of their uses seem to fit the description perfectly, they often enough do not converge on creating a unified “unilateral” picture. Specifically, I will argue that *most* does not manifest an absolute “unilateral” pattern, and that the numbers do not manifest a perfect “bilateral” pattern. Rather, depending on the appropriate ad hoc interpretation they receive in context and the properties associated with that interpretation, *most* and the numbers demonstrate either a “unilateral” or a “bilateral” pattern.

Section 4.1 will present the received view picture, according to which *most*, but not the numbers, seems to display the “unilateral” pattern. In section 4.2 I will demonstrate that *most* does not always go the “unilateral” way, and that similarly, the numbers do not uniformly follow the “bilateral” pattern. I will then try to account for the differential patterns of both types of expressions by reference to three parameters: punctuality, interactional distinctness from higher values, and the enablement of part-whole inferences (4.3). It is the value on these parameters which determines whether a given expression will manifest some “unilateral” or “bilateral” pattern of behavior. If so, the observed differences between *most* and the numbers do not constitute a valid basis for rejecting my analysis of *most* as circumbounded.

4.1 A unilateral pattern for *most* but not for the numbers

The examples in 4.1 are all cases where *most* is interpreted differently from the numbers, arguably justifying a unilateral analysis only for *most*. First, note the following from Geurts (1998:106):

- (21) a. ~At least/at most/half of/exactly **two hundred** ships.

- b. ~?At least/at most/half of/exactly **most** of the ships.

It is certainly true that the scalars are not as free as the numbers to combine with the modifiers in (21). Indeed, a search of BNC and CoBuild showed a meager number of *at least mosts*, in most of which *at least* did not have scope over *most* alone. I found no *at most most*, nor *exactly most*. Second, Sadock (1984) argues that the cardinal scale, but not the quantifier one, is easily reversible. Indeed, all 26 *most if not* occurrences in the BNC were followed by *all* – e.g., (23) (rather than by *some*). Here, however, is an example with a reversed scale for a numeral:

- (22) After **two** hours, if not sooner (Collins CoBuild).

Next, we come to more pertinent cases, where the difference between the numbers and *most* (and other scalar quantifiers) seems to directly motivate a unilateral analysis for *most*. First, the ability to shift smoothly from *most* to *all* in the same utterance, without sounding contradictory has been cited in support of a unilateral semantic analysis, because it shows *most*'s compatibility with 'all':

- (23) Right I'm gonna ask the director, who I think is going to pick up **most if not all** of your points, <BNC: J43 182>

Second, *most* is predicted (by Fred Landman, p.c.) to pattern with *at least n* expressions as a discourse antecedent for *they*. The numbers, on the other hand, can only be interpreted as 'exactly n' when they serve as discourse antecedents (Kadmon 1987). Consider the following examples, assuming that that the facts are that all the Hondas (say, 20) were defective:

- (24) a. ~**At least 11** Hondas were defective. They took **them** out of the shop.
 b. ~**Most** Hondas were defective. They took **them** out of the shop.
 c. ~**11** Hondas were defective. They took **them** out of the shop.

Kadmon predicts that (24a) is acceptable and true with the interpretation that 'all the Hondas were taken out of the shop'. (24c), on the other hand, is infelicitous, because the antecedent specifies 'exactly 11', and the known facts are that '20 Hondas are defective'. Landman proposes that (24b) patterns with (24a), rather than with (24c), arguing that since *most* is compatible with 'all', if all the Hondas were defective, then all of them were denoted by *most*, and therefore 'all' is denoted by *them*.

(19) above is a question from my *most* questionnaire, showing that most of my subjects interpreted a discourse anaphoric *they* as referring to 'possibly all', when the antecedent was *most*. Indeed, 31/36 (86.1%), similarly chose the same answer for *more than half*. It therefore seems that Landman's prediction for a lower-bounded-only pattern for *most* is confirmed by my subjects.

Last, behavior under negation has also been used to argue for the differential behavior of the scalar quantifiers and the numerals. The unilateral received view seems to have a straightforward explanation for the observation that negating scalar expressions such as *most* is normally taken as negating only their lower bound. If *most* covers any quantity above 50%, 100% included, then negating it can only be compatible with less than 51%. Indeed, *not most* in the following is probably interpreted as 'sometimes' i.e., as less than the lower bound (51%), and not as 'always', i.e., as more than the upper bound (99%):

- (25) A: Uh, well that's true. But I get my hair cut about what every six weeks too.
 B: Yeah, and it looks bad most of the time too. Well **not most** of the time. (LSAC).²²

Since 100% is assumed to be a possible denotation for *most*, the received view can also straightforwardly account for the observation (by Horn 1996) that while the numbers do not allow confirmation of a lower value (by an initial *yes*) if the reality is that a higher value is the case, the scalar quantifiers do (see also the results reported on in Papafragou & Musolino 2003). Here are Horn's examples in this connection (his examples 22, 22'), and the reader can substitute *most* for *many*:

- (26) a. ~A: Do you have **two** children?
 B1: **No**, three.
 B2: ?**Yes**, (in fact) three.
 b. ~A: Are **many** of your friends linguists?
 B1: ?**No**, all of them.
 B2: **Yes**, (in fact) all of them.

Note that in both cases B corrects A's lower value (*two*, *many*) to a higher value (*three*, *all*). The difference is that the number correction takes an initial *no*, whereas the quantifier correction takes an initial *yes*. Presumably, this is because '*three*' is not included in the denotation of *two*, but '*all*' is included in the denotation of *many*. The same should apply to *most*. In sum, it seems that the numbers and the scalar quantifiers manifest different interpretative patterns. The numbers go the "bilateral" way, the quantifiers go the "unilateral" way. Such findings seem to pose a problem for my circumscribed analysis of *most*.

4.2 Inconsistent "unilateral"/"bilateral" behaviors

We now consider a different set of examples, ones where *most* patterns as a bilateral expression, and the numbers pattern as unilateral expressions. The findings demonstrate that the differences between *most* and the numbers are not as consistent as they are presented in the literature. While the distributional pattern identified by Geurts (1998) (see (21) above) is on the whole supported by the data I examined, I did find two *almost mosts* on the web (as well as many *almost a majority*), a few *at least mosts* (see again (5)), and 4 *roughly mosts*. And here's a constructed example where I think that *exactly* would be appropriate:

- (27) ~A: It sounds like your favorite authors are Castel-Bloom and Shabtai. Have you read all of their books?
 B: I've read 7 by Castel-Bloom. I think that's most of her books. But I've only read two by Shabtai. I doubt that's most of his books.
 A: That's **exactly most**. He only wrote three books.

When the total set is three books, *most* necessarily refers to one specific value, namely 'two', under my analysis. Hence the ability to modify it with *exactly*.²³

²²All the examples from LSAC come from research conducted together with John Du Bois.

²³While audiences exposed to this example seemed to accept my judgment, Laurence Horn (p.c.) and one of my referees find it an impossible combination. As we shall see later, even if they are right, this has no bearing on the upper bound issue.

Next, recall Sadock’s (1984) point about the non reversibility of the quantifier scale, as opposed to the numeral scale, which is reversible. First, note that scalar comparisons tend to go from lower to higher values not just for *most*, but for the numbers as well: All 4 *two if not* were followed by higher values in the BNC. The 3 numeral + *if not* combinations in LSAC were also all followed by higher values. The 2 cases where a numeral + *if not* was followed by a lower value in CoBuild referred to time (see (22)). The other 2 cases specified a higher value. Thus, in general, all scales are not often reversed. Moreover, the following set of examples shows how the quantifier scale can in fact be reversed. Koenig (1991), justifying a bilateral analysis for the numbers, argues that measure phrases show context-dependency as to what the direction of the scale is. The following examples have been constructed on the basis of similar examples from Koenig with numbers (Koenig stars the counterparts of examples (b) and (c) with numbers). As can be seen, a similar pattern is exhibited by the scalar quantifiers:

- (28) a. ~You can buy this book with **all/most** of the money you have there, in fact with **most/some** of your money.
 b. ~??You must pay for this book with **all/most** of the money you have there, in fact with **most/some** of your money.
 c. ~??You can buy this book with **most/some** of the money you have there, in fact with **all/most** of your money.
 d. ~You must pay for this book with **some/most** of the money you have there, in fact with **most/all** of your money.

Thus, it may just be the case that there are less uses for “reversed” quantifier scales. But they are not irreversible in principle.

Next, I quote examples where the pattern adduced in support of the unilateral meaning of *most* is equally attested for the numbers, as well as for nonconventional scalar items:

- (29) a. Well yeah but you see that the trouble is they’ve been now **two if not three** pilot phases. <BNC: H5E 929>
 b. her services if not **in regular demand** are no doubt **in demand**.
 <LLC: 12 4b 3 6870 1 1 a 11 2>
 c. Unfortunately for such critics it has been found that acupuncture works **equally well, if not more effectively**, on animals. <BNC: CB9 1459>
 d. PHIL: ... And they’re **very close together**,
 ... in fact ... they’re in a v- ... **joined uniform state**. <SBC: 027>²⁴

It’s quite clear that upward compatibility, enabling the noncontradictory shift from *most* to *all* is not unique by any means. (29a) shows it for a number, and (29b)–(29d) are examples where a noncontradicting scalar upgrade is available for predicates that do not at all form a conventional Horn scale. The scalar quantifiers do not pattern differently from the numbers, then, nor for that matter from nonconventional scalar predicates.

²⁴Impressionistically, however, it seems that the numbers and nonconventional scalar predicates more often than the scalar quantifiers are modified by *at least* when followed by an *if not*. This stems from the range/punctuality difference, however (see 4.3).

A central argument of Koenig (1991) for preferring a bilateral analysis for the numbers is that their scalar effects are not stable. This undermines the usefulness of the conventional scale as a **linguistic** tool, a necessary feature of the received view. The examples below demonstrate that the same is true for *most*. Thus, just as Koenig notes that downward entailments, the basis for a Horn Scale, are not always applicable to the numbers, they are not always applicable to the scalar quantifiers. Just like (30a) does not entail (30b), so (31a) does not entail (31b) (both receive a collective reading):

- (30) a. Usually, **four** men carried a corpse; (Collins CoBuild)
 b. Usually, **three** men carried a corpse.
- (31) a. PATTY: ... that was the last time they were **a=ll** together. <SBC: 023>
 b. That was the last time **most/some** of them were together.

Koenig further notes that downward entailments are not guaranteed even when the numbers are given distributive readings as in:²⁵

- (32) JIM: ... if we had [...] **ten** accounts,
 FRED: [Okay].
 JIM: ... we would charge (H) ... five-hundred fifty dollars on ac- on an account,
 <SBC: 014>

which does not entail that 'If we had nine accounts we would charge ...'. The same is true for *most*. *~If we had most accounts, we would charge ...* does not entail 'If we had some accounts, we would charge ...'.

Next, Koenig discusses examples similar to (33), in order to show that number measure phrases too are not downward entailing:

- (33) MELISSA: they live **three** blocks away, <SBC: 019>

Obviously, (33) does not entail that 'they live two blocks away'. Similarly, *all the time* in (34) does not entail 'most of the time':

- (34) a: let me give you [@m] five hundred pounds or something
 A: yes
 a: and instead of you ringing me up **all the time**. <LLC: 22a1521035012a20>.

Since scales depend on a downward entailment relationship, and since these entailments are not always applicable, scalar effects cannot be lexically/automatically determined. This is true for both the numbers and the scalar quantifiers, as we already saw in (18) above.

Finally, in arguing for a 'set of exactly n' analysis for the numbers, Kadmon (2001:69) notes another nonstable scalar effect. Numbers can receive an 'at least' reading only in argument positions. Predicative positions are restricted to the circumbounded interpretation. She notes that in an example such as (35), *three* cannot be interpreted as 'at least three', so that the proposition is false if there were four cats:

²⁵See Chierchia (2004) for a structural explanation for these instabilities (by reference to *any*-licensing contexts). The crucial point for my argument, however, is the similar patterning of *most* and the numbers. I assume that the same explanation applied to the numbers can apply to *most* as well.

- (35) Why, they were **three** cats! (Lewis Carrol, *The three cats*. A story excerpted from a letter by Carrol to some of his child friends)

The same is true for *most*. Note that according to Saddam Hussein, (36) must be false, given that he was elected president with a hundred percent of the votes (as reported on Kol Israel radio, 10.18.2002):

- (36) ~The votes for Saddam Hussein were **most** of the votes cast in the election.

Thus, *most* too shows an asymmetry in receiving lower-bounded-only readings.

Let us now re-examine Landman's (p.c.) proposal that *most* can naturally serve as a discourse antecedent for a *they* interpreted as 'at least most'. We saw that this was the case in the questionnaire question cited as (19) above, but note that quite similar results were obtained when I asked about *61* in the same question:

- (37) It is not necessary for all 120 Knesset members to be present for the Knesset to convene. The law requires that **61** members participate in the assembly discussion tomorrow. **They** will be asked to vote on a series of social laws.

Question: Who are the "they" who "will be asked to vote on a series of social laws"?

Answers: **A.** All the Knesset members who will participate in the assembly discussion tomorrow, that is, no less than 61 and up to 120 Knesset members, including 120 Knesset members.

B. All the Knesset members who will participate in the assembly discussion tomorrow, that is, no less and no more than 61 Knesset members.

Recall that a large majority interpreted *most* as 'possibly all' in the counterpart question. Similarly, 17/19 (89.5%) chose an 'at least 61' interpretation for the pronoun in the *61* case too. In fact, there were more subjects (16.7%) who insisted that the *most* anaphoric *they* only refers to an upper-bounded 'most but not all', than on an upper-bounded 'exactly 61' answer for the *61* question (2/19, 10.5%). A comparison between *most* and *61* as antecedents for *they* then shows that they pattern quite similarly ("unilaterally"), the 'at least' reading being preferred for both in this case. Similarly, recall that Horn (2005) predicts that whoever bets *most* wins the bet if 'all' is true (see again (17)). For a counterpart number, however, he is not sure whether the person betting on a specific number wins when a higher number turns out to be the case. In fact, I think that he should predict that once the bilateral number is incorrect the better should lose. In any case, as we saw, *most* and 80% patterned almost identically in the very similar guessing game question (see (18) above).

However, in the questionnaire reported on above I actually created two types of contexts where *most* or a number were later referred to by *they*. In both, subjects were told that reality is such that the predicate of the antecedent clause is true for 'all' (for *most*) or a for higher number (for the number antecedent). But only in one type of context (19) were the subjects willing to say that *most*'s conveyed meaning was 'at least most'. Hence, so was the anaphor. In the other type of context, subjects were not ready to assume that the speaker intended them to interpret *most* as conveying 'possibly all', despite the fact that 'all' was true in reality. Rather, they chose to see a gap between reality and the explicit utterance, which they interpreted according to its lexical,

unenriched circumbounded meaning. Hence, so was the anaphor interpreted as well. This is clearly contra Landman's claim. If *most* can denote 'all' and we know that 'all' is the case, the pronoun should have referred to 'all'. But it didn't in most cases. Similar questions with a number antecedent produced similar results (the reader is referred to Ariel 2003 for a more detailed discussion of the questionnaire results). What this means is that when subjects had to interpret *they* on the basis of the linguistic meaning of *most* or a number, they tended to go for a circumbounded interpretation. In other words, given the same contextual assumptions, *most* and the numbers patterned in the same, "bilateral" way.

Finally, I would like to show that the assumption that negating *most* always amounts to denying only its lower (51%) boundary cannot be maintained. First, a search for *not* + 0–3 words + *two/hundred* (I picked one small and one larger number) in LSAC revealed that for the numbers too negation is mostly interpreted as 'less than the specified number' (23 cases), rather than as 'more than the specified number' (3 cases), a 7.66 gap.²⁶ Next, confirming a higher number does not always require an initial *no* (this is then a "unilateral pattern" for a number). In the following, it seems that *30* patterns with the scalar quantifiers rather than with the numbers. Compare (26a) with (38), where the number (*30*) patterns with the scalar quantifiers, prompting an initial *yes*, rather than *no*:

(38) ~A: Do you have \$**30**?

B1: ?**No**, \$**32**.

B2: **Yes**, \$**32**.

The same seems to be true in the following cases:

- (39) a. MONTROYA: (H) Italy for example,
 ha=s ... a voter ... turnout,
 ... of **ninety** percent. <SBC: 012>
 ~B1: **Yes, ninety three** percent.
 ~B2: ?No, **ninety three** percent.
- b. MONTROYA: There were **three hundred thousand** people.
 ... at that march in
 nineteen-sixty-three. <SBC: 012>
 B1: **Yes, three hundred and ten thousand** people.
 B2: ?No, **three hundred and ten thousand** people.

In fact, versions similar to B2's response in (26a) have always been considered acceptable in the literature, when considered in the "social security" context (where it is assumed that one gets government support if one has at least two children). And compare (26b) with (40), discussing the possible extinction of the California Condor:

(40) (Imaginary Conversation in 1987, when there were only 32 California Condors in existence).

~A: Are **most** of the California Condors infertile?

B: (Sadly) **No, all** of them are.

²⁶The count for *a hundred per cent*, was not included. Of course, all such cases were interpreted as 'less than'.

In (41a) the speaker corrects himself from *most* to *every*. Similarly in (41b). Note that the corrections are prefaced by the “bilateral” *no*, rather than by the “unilateral” *in fact*. (41c), on the other hand, shows a “unilateral” *yes* response despite the denial of the lower (and not upper) bound on the scalar *beautiful*:

- (41) a. In my thirty years of ??,
we have lost 1200 football games,
600 basketball games,
most of our –
No, EVERY track meet. (“The Ellen show”, 8.22.2003).
- b. R.G: How are decisions made, according to **(the) majority**?
O.L: **No, never**.
R.G: So who decides?
O.L: We **ALL** do (Originally Hebrew, University committee meeting, 9.1.2004).
- c. Emma: e-that Pat isn’t she a doll?
Margy: Yeh isn’t she pretty
Emma: Oh: she’s a **beautiful** girl.
Margy: **Yeh** I think she’s a **pretty** girl
(Heritage 2002:ex. 24, transcription simplified).

(41) then show that contra current assumptions, initial *no* can be appropriate for a *most* to *all* correction (an upper boundary denial), and an initial *yes* may be appropriate for a *beautiful* (a strong scalar term) to *pretty* (a weaker scalar term) correction (a lower boundary denial).

All in all, the numbers sometimes manifest the “unilateral pattern” ((26a) sometimes, (29a), (37), (38), (39)), and the scalar quantifiers sometimes manifest the “bilateral pattern” ((27), (28), (31), (34), (36), (40), (41a,b)). The differences between them cannot therefore be absolute. Now, contrast the nonstable judgments in (26) with the following, where judgments concerning the initial *no* are not easily manipulated:

- (42) ~A: **Most** condors are infertile.
B1: ??**No**, you’re wrong. **80%/99%** of them are.
B2: **No**, you’re wrong. **All** of them are.

Since 80%/99% fall within the range covered by *most*, B1’s *no, you’re wrong* creates a contradiction (unless interpreted metalinguistically). It is not, however, necessary to interpret B2’s *no, you’re wrong* metalinguistically in order for it to be an acceptable response. The (nonmetalinguistic) judgments in (42) are stable, attesting to the semantic nature of the problem in B1’s utterance. This is not the case for the previous examples we discussed in 4.2.

Finally, note that even the various scalar quantifiers don’t always pattern in the same “unilateral” manner. *Half*, for example, patterns with the numbers regarding the cooccurrence restrictions noted in (21), rather than with the scalar quantifiers:

- (43) a. **Exactly half** of my life has been here. (LSAC)
b. I think we need to get everybody back together for a **at least half** an hour. (LSAC)

In fact, it's even possible that *some* and *most* do not manifest exactly the same distributional restrictions either. I couldn't find any *at least most* cases in any of the corpora searched except for the web, but I did find such *some* combinations in LSAC (11 cases, but note that *some* is by far more frequent than *most*). Note also that substituting *some* for *many* in (26b) improves B1's response dramatically:

(44) ~A: Are **some** of your friends linguists?

B1: **No**, all of them.

We conclude that it's not the case that *most* is invariably interpreted unilaterally. Similarly, it's not the case that the numbers are invariably interpreted "bilaterally". In addition, we noted differences between the scalar quantifiers themselves (*half* vs. *most*, *some* vs. *many/most*). In 4.3 I will venture an initial proposal regarding the differences between *most* and the numbers. I will argue that they follow naturally from prototypical, rather than absolute differences between the numbers and the scalar quantifiers. Crucially, these differences are orthogonal to the question of upper bound.

4.3 Deconstructing the "unilateral" pattern

The received view is right about *most* and the numbers manifesting different interpretative patterns (4.1), but I claim that these are only statistical tendencies: In 4.2 we have reviewed cases where *most* patterns according to the "bilateral pattern", and where the numbers pattern according to the "unilateral pattern". In addition, while both *some* and *most* often pattern in a similar ("unilateral") way, they do not always pattern uniformly. I would like to propose below that the so-called unilateral pattern actually consists of a number of patterns, which are in principle independent of each other. Crucially, they derive from factors which are irrelevant to the presence/absence of an upper bound: (Non)punctuality, degree of interactional distinctness and (non)enablement of part-whole inferences.

I will argue that what is termed a bilateral behavior comes about when the expression is interpreted punctually and/or as interactionally different from higher values not denoted by it, and/or when part-whole inferences are blocked. What is considered a unilateral behavior comes about when the denotation is perceived to cover a wide range, and/or not very distinct from higher values, and/or when part-whole inferences are encouraged. Statistically, numbers frequently meet both former conditions. At least, they are punctual and distinct from higher values. This is why they are considered "bilateral". *Most*, on the other hand, tends to meet the latter conditions (it is nonpunctual, interactionally not so distinct from the higher 'all' value, and perhaps enables part-whole inferences more often). This is why it is considered unilateral. My argument is, first, that the interpretative pattern results from these parameters directly, rather than from the fact that some expression is a number, and therefore bilateral, and another is a quantifier, and therefore unilateral. My second and main point is that what is taken as the unilateral pattern does not actually entail lack of upper bound for the expression. The differences below are orthogonal to the upper bound issue.

We start with (non)punctuality. Clearly, scalar quantifiers such as *most* denote a wide range, while the numbers' meanings are punctual. Indeed, in Ariel (2002), I presented a statistics based on two SBC conversations that 84.5% of the numbers are interpreted as 'exactly n'. Only 15.5% receive an 'about n' reading. I propose that this difference accounts for why the scalars are not

so much in need of modifiers such as *at least*, which create a range out of a punctual value. Thus, the low frequency of the combinations in (21b) is due to the fact that the numbers are punctual, whereas *most* covers a wide range. Note that *half* too is punctual. If I am right, it should therefore pattern with the numbers rather than with the scalar quantifiers with respect to cooccurrence with *at least* and *exactly*. Indeed it does (see (43) again). The punctual *half of* is also inappropriate for a nonpunctual quantity such as ‘most’. Papafragou (2003) similarly finds a difference between *half* and aspectual verbs, such as *start*, which are vaguer in terms of their boundaries in triggering upper-bounded interpretations, and argues for a “discrete/nondiscrete asymmetry”. In fact, as we have seen, once *most* receives a punctual interpretation (in (27)), it too can be modified by *exactly*. Be that as it may, I fail to see the connection between being wide range and being only lower-bounded. A wide range expression may very well be upper-bounded too. Whereas *at least/more than half* are wide-range lower-bounded-only expressions, *between twenty and a hundred* (LSAC) is wide-range, but upper bounded (see Keenan 1996), as are *a minority*, and *at most x*. Thus, whether or not *most* carries a lexical upper bound is quite orthogonal to the statistical tendencies noted by Geurts (1998). These reflect the fact that numbers tend to be interpreted punctually, and most of the scalar quantifiers (but less so *half*) tend to be interpreted as covering a wide range.²⁷

Another difference between *most* and the numbers is that any number has an infinite number of higher values. *Most* does not. First, objectively speaking, the difference between 99.99% (the highest value covered by *most* under my analysis) and 100% is strikingly smaller than the difference between any number, as high as it may be, and other higher values not coded by that number. Second, while the difference between *most* and *all*, despite its small objective size, has been taken to be of utmost importance in the literature, interactionally, speakers and addressees do not necessarily assign the advantage of 100% over 99% such an important role (see Ariel 2004). Most probably because we have very low expectations to be making universal claims, the difference between ‘most’ and ‘all’ is not usually so crucial interactionally. Either type of generalization is seen as strong evidence/justification for some conclusion/course of action (see (1) again). In other words, I claim that *most* is not interactionally very distinct from *all*, because in most contexts (see Ariel 2004), the difference between ‘most’ and ‘all’ is not relevant, in that ‘all’ does not carry significantly different contextual implications that ‘most’ doesn’t, to use Sperber and Wilson’s (1986/1995) terminology. In addition, while a universal generalization is in some sense stronger, it actually brings with it a great vulnerability, which drastically weakens it. One counter-example is sufficient to render the universal proposition false. It is much harder to falsify a *most* claim, especially since in natural conversations, the size of the whole set is quite often not obvious to the interlocutors (see (1), (5)). In other words, the slight advantage of *all* over *most* in strength is offset by a large disadvantage in refutability. The following example testifies to this:

(45) A: Why do you **always** think you’re right?

B: Because I’m RIGHT **most** of the time
(Beetle Bailey, *Int’l Herald Tribune*, 10.17.2003).

Note that B accepts A’s proposition that he ‘always thinks he’s right’, but he doesn’t support it with a universal claim. A *most* claim is strong enough. Justifications by *most* are often not

²⁷Of course, punctual values too, even if modified by *exactly*, have some ‘slack’, to use Lasersohn’s (1999) term.

interactionally very different from justifications based on *all*. If I am right, then *some* and *most* should differ on this point, because unlike *most* and *all*, the (objective) difference between *some* and *all* is significant. A *some* generalization plays a different (and weaker) interactional role from that of *most* (and *all*), no doubt because *some* doesn't even have to denote a majority. I therefore expect *some* to sometimes pattern with the numbers and not with *many/most*. Indeed, we saw above that it was not so difficult to find *at least some* examples (but since *some* too covers a wide range, it's not surprising that there were no cases of *exactly some* in LSAC). Also, an initial *no* when confirming 'all', I claimed, was more natural for negating *some* than *many* and *most* (see (26b) vs. (44)). In fact, the difference between *some* and *most* seems to be more significant interactionally than that between *most* and *all*: 4/14 (28.6%) of the corrections of *some* to *most* were explicitly marked as a correction to a different alternative, but this was so for only 1/22 (4.5%) of the corrections from *most* to *all* in LSAC. For *most/all* alternations, the speaker seemed to merely be wavering between the two options (see 6.1 in Ariel 2004).

The result of the differences in punctuality and interactional distinctness from higher values between most uses of the numbers and *most* is that the numbers tend to contribute towards a proposition which is more specific and more easily distinct from another asserting a higher number, but under most circumstances this is not the case for *most*. Denying it is then expected not to be prefaced by an explicit *no*. If this is so, we can understand why researchers were led to believe that the difference between the numbers and *most* concerns the upper bound. Positing an upper bound only for the numbers seems to explain the different pattern of interpretation for *most* and the numbers under negation, as well as the differential interpretation they seem to receive as discourse antecedents. This, however, is problematic in that it would predict an absolute difference, which we don't find.

To see the significance of (non)explicitly stated distinct alternatives, we can examine the behavior of numbers and *most* under negation. I claim that the degree to which some value is interactionally distinct from another, higher one which a responder is proposing, interacts with the role of negation to produce the interpretative patterns observed in the literature. As we have seen for the numbers in 4.2, normally, negation is pragmatically interpreted as suggesting that a weaker concept is the case. Indeed, based on a variety of psycholinguistic experiments, Giora et al. (2004) and Giora et al. (2005) argue that negation is often pragmatically used as a hedge. The negated element receives a mitigated interpretation of 'less than' (note the ambiguity of *less* as 'lower in quantity' and as 'lacking', as in *merciless*). For *most*, less than 'most' means less than 51%, and not 100%, which is more than 99%. Horn might counter that negation seems to be a mitigator just because the negated elements here are scalar, and only have lower boundaries.²⁸ But this is equally true for nonscalar terms. *A fork*, for example, is obviously a circumscribed concept. Still, one does not normally deny *a fork* in order to affirm 'more than a fork':

(46) ALINA: and I didn't need a fork <SBC: 006>

We routinely interpret *not a fork* as 'less than a fork', say, 'no fork', and not as 'more than a fork', say, 'a fork and a knife'.

However, despite the fact that statistically, negation tends to be interpreted as applying to the lower bound of all concepts, there is an intuition about a difference between the numbers and the scalar quantifiers (see again Horn's examples in (26)). I suggest that once we take into consideration the examples in (38), (39), (40), (41), (42), (44), where the pattern is reversed for

²⁸See, however, the discussion in 3.2.

both *most* and the numbers, it becomes clear that something else must be at work here. The generalization that emerges from all the examples is that the decision between an initial *yes* or *no* depends on the importance attributed to the difference between the two values in the given context. Large differences tend to be prefaced by *no*, small/insignificant differences tend to be prefaced by *yes*. As already mentioned above, since substituting *some* for *many* in (26b) improves it, it corroborates my claim that what matters is that there be a relevant difference between the answer and the original alternative suggested in the question. The following surprising denial attests that indeed, *no*'s have more to do with interactionally significant differences than with semantic denials:

- (47) A.S: Do you have roasted **almonds**?
 Salesman: **No**.
 A.S: And what are these, aren't these almonds (pointing to roasted almonds)?
 Salesman: These are **halved almonds** (Originally Hebrew, 3.10.2004).

In this case, the salesman thought that whole and halved almonds constitute significantly different alternatives.

What the answers in (26), (38), (39)–(42), (44), (47) show is that the decision to answer positively or negatively does not (directly) hinge on the number-quantifier contrast. Rather, *yes* and *no* seem to mark different stance alignments (to use John Du Bois' term, p.c.) between speakers: *Yes* marks agreement, *no* marks disagreement. These stances depend of course on the objective content conveyed by the interlocutors. A higher degree of agreement on content tends to trigger *yes*, and a low degree of agreement tends to trigger *no*. But what count as large vs. small differences are not necessarily objectively measured. Specifically, both *no* and *yes* can be used when speakers disagree on the facts, as is the case in virtually all the examples we reviewed in section 4. It seems that where the difference between 'two' and 'three' ((26a) uttered in most contexts), between 'many' and 'all' (in (40)) (extinction will only follow once all the birds are infertile), and between *most* and *every* (in (41a)) matters to the speaker, supplying a higher number/quantity justifies an initial *no* (see also Carston 1990). Where the difference between 'many' and 'all' (in (26b), between the numbers in (38), (39)), and between 'beautiful' and 'pretty' is less crucial (Heritage 2002:218 says that Margy in (41b) creates a "pseudo-agreement"), an initial *yes* is preferred.²⁹ Given that the choice between initial *yes* and *no* is interactionally dictated according to how relevant the difference between the different positions is taken to be, we can account for the intuition that *most-to-all* corrections tend to trigger an initial *yes*, whereas a number-to-a higher number correction tends to trigger an initial *no*. As I argued above, the latter difference is more often objectively and interactionally more relevant than the former.

A related factor at work here, contributing to the preference for an initial *no* for the number corrections, but not for *most*, is again a function of the tendency to interpret numbers punctually,

²⁹Interestingly, Horn (2005) draws upon a similar distinction in his account for why (i) is marginally acceptable, but (ii) is not:

- (i) ~A: Is dinner almost ready?
 B: ?Yes, it's ready (Horn's (24a)).
- (ii) ~A: Is Fredo almost dead?
 B: #Yes, (in fact) he's totally dead (Horn's (24b)).

There is a smaller gap between dinner being ready and dinner almost being ready than between Fredo being totally dead and Fredo being almost dead.

whereas *most* as covering a wide range. Just because the numbers denote precise values, it is plausible that the speaker intends an ‘exactly’ enrichment when she uses a number. Since she chose such specific quantification, it is likely that she meant no more and no less than that value. Thus, the common interpretation of the numbers as ‘exactly n’ accounts for the relative ease of denying a lower number when a higher one is the case. It would follow that speakers would perceive as more crucial a correction of a punctual value than of a wide-range value. Hence the intuition that a corrected number value cooccurs with an initial *no* more than a scalar quantifier. If I am right, and the difference pertains to punctuality vs. range, rather than to numbers vs. *most*, then when numbers are interpreted as ‘about n’ (round numbers normally – see Jespersen 1949:586–7), they should pattern with *most*, rather than with nonround numbers, so they too should show a preference for an initial *yes*, rather than *no*, when the responder proposes that a higher number is the case. This is indeed so, as we have seen in (39). *Half*, I predict, should be a hybrid case. While it is punctual, it seems to have a rather large “slack” (see the results in Papafragou & Schwarz to appear), larger than nonround numbers, at least. The distributional differences between *most* and the numbers re negation can then be attributed to the punctual-range difference between them, and to the difference in whether they constitute contextually distinct alternatives, combined with stance (dis)alignment decisions. Since the numbers are punctual, correcting them in any direction more easily justifies a disagreeing stance (indicated by an initial *no*). Since the difference between *most* and *all* is not very crucial interactionally in many cases, speakers don’t usually feel they have to adopt a disagreeing stance when correcting a *most* to *all*.

A third factor involved in the occurrence of a “unilateral” interpretative pattern is the applicability of downwards entailments and upward compatibility, or as I call them, part-whole inferences. As we have already seen in section 3, this is a pragmatic phenomenon. Downward entailments, argues Koenig, following Anscombe & Ducrot (1983), can be inferred based on our world knowledge. It is our world knowledge that tells us that if *Vaska and I drank four bottles of brandy between us* (Collins CoBuild) then ‘Vaska and I must have drunk three bottles of brandy between us’, but “such entailments derive from our world-knowledge of consumption” (Koenig 1991:145). The same is true for *all* and *most/some*. Sometimes the entailment is accepted, at other times, it is not. Newstead (1995), for example, found that a great majority of his subjects said that ‘some’ is not implied by ‘all’. Larger parts (‘all’, ‘most’) often, but not invariably, entail that smaller parts (‘some’) are true (see (18) above). Similarly, smaller parts are sometimes, but not always, seen as compatible with the predicate applying to larger parts/wholes. We therefore do not need specialized lexical scales for this purpose.

To see that upper boundedness is only statistically correlated with blocking of part-whole inferences, consider the following:

(48) Vaska and I drank **four** bottles of brandy between us. (Collins CoBuild)

(48) seems to be true in case ‘we drank five bottles ...’, and allows for an inference that ‘we drank three bottles ...’. Similarly,

(49) **Most** Israelis decided for peace.

seems to be true in case ‘all Israelis decided ...’, and allows for an inference that ‘half/some of the Israelis decided ...’. In these examples we see that both the numbers and *most* enable upward compatibility and downward entailments (but see the questionnaire results for (18)). Now, what about a bounding expression such as *between 70% and 80%*? Its compositional

meaning restricts it to ‘no less than 70% and no more than 80%’ (see Keenan 1996), which should be incompatible with states of affairs in which ‘90%’ is true. Unlike *most*, bounding expressions explicitly exclude the application of the predicate to the higher values. Even so, 23/23 of my students chose an ‘at least 70–80%’ interpretation for the counterpart expression in Hebrew when I created an appropriately biasing context:

(50) Maya is a new teacher, who will start teaching at a school next year. She very much wants to do well on her job. Galit, the principal, tells her that she has decided to assign her to an especially weak class. The following conversation takes place at the beginning of the school year:

- Galit: The Ministry of Education tests will be conducted in 6 months. It’s a good idea to start preparing the students, especially your class. If half of the students in your class pass the tests it’ll be an excellent achievement.
- Maya: I promise you that **between 70% and 80%** of the students will pass the tests.
- Galit: Let’s hope so. You know that the education system has no provision for bonuses, but if this miracle happens, I will get you a 3,000 sheqel bonus.

7 months later it turns out that 90% of Maya’s students passed the tests.

QUESTION: In your opinion, will Maya receive the 3,000 sheqel bonus?

100% of the subjects chose “yes”. Thus, even a noncontroversially upper-bounded expression such as *between 70% and 80%* can be taken by subjects as compatible with a higher (90%) value, given the right circumstances. The lexical upper bound seems unable to block part-whole inferences here when they are called for contextually. The (linguistic) question of upper boundness is therefore independent of the (pragmatic) question of whether part-whole inferences are allowed.

Note that just like “wide-range” numbers take the “unilateral pattern”, so should a punctual quantifier, such as *half*, take the “bilateral pattern”. Indeed, Papafragou & Schwarz (to appear) show that subjects refused to accept ‘all’ as compatible with *half*. In other words, they refused to allow for a part-whole inference here. *Half* here patterns with the numbers, rather than with *most* (in fact, it seems even more “bilateral” than the numbers). Part-whole inferences tend to be blocked or encouraged according to contextual circumstances (This is my interpretation of the results reported in Papafragou & Schwarz to appear).

All in all, section 4 has argued that *most* patterns with the numbers in that the scales they participate in are subject to the same pragmatic constraints. Whether or not downward entailment works, whether or not the scales are reversible, their behavior under negation and their behavior as discourse antecedents are all pragmatically determined, and hence, show statistical tendencies, rather than absolutely different patterns.³⁰ These differences, however, are due to the punctuality and interactional distinctness differences, as well as to the (non)applicability of part-whole inferences, all of which are irrelevant to the question of the upper bound. In other words, I am proposing that the differences noted in the literature are not defined over numbers

³⁰Alternatively, if Chierchia’s (2004) account is applicable to all the cases above, then the factors involved are semantic rather than pragmatic.

vs. scalar quantifiers per se. The “Bilateral pattern” is statistically associated with the numbers more often than with the scalar quantifiers, because the numbers tend to be punctual, and often constitute very well-defined distinct alternatives. The scalar quantifiers tend to manifest a “unilateral pattern”, because they tend to denote a wide range, and at least for *most*, the difference between the quantifier and the higher alternative (‘all’) is not usually interactionally relevant. The nature of the exceptions to the prototypical cases proves that it is the parameters mentioned above that are responsible for the interpretative pattern associated with some form. This is why round numbers (nonpunctual) resemble *most*, why punctual *most* (when the whole is 3 or 4) resembles the numbers, why *some* (clearly more distinct than *most* from the higher value alternative ‘all’) resembles the numbers, and why *half* (often punctual) resembles the numbers. In addition, part-whole inferences are performed according to our world knowledge, and applied if contextually appropriate. These are equally relevant for *most* and the numbers. I am not even sure that part-whole inferences are more often appropriate when scalar quantifiers are involved. This is why downward entailments don’t always apply to both numbers and to scalar quantifiers. It remains to be seen whether the history of the analysis of the meaning of the numbers will not repeat itself for quantifiers and other scalar predicates. Be that as it may, the facts noted in section 4 demonstrate that the different patterns sometimes manifested by *most* and the numbers do not pose an obstacle to analyzing *most* as lexically upper-bounded, just like the numbers are.

5 No lower bound?

Thus far, we have concentrated on the question of upper bound for *most*, ignoring *most*’s lower bound. Indeed, there seems to be no controversy in the literature about the lower bound for *most* (nor for *majority*). Note, however, the following examples, where Hebrew *rov* ‘majority’ and English *most* denote less than 51%, violating the lower bound assumed by the received view, as well as by my analysis:

- (51) a. Yosifov was elected by a **majority** of 41.6% ...and does not need a second round (Originally Hebrew, *Haaretz* 6.5.2003).
- b. Knesset Member Shimon Peres was elected yesterday as temporary chair of the Labor party ... by a **majority of less than 50%** (49.2%; 631 votes) (Originally Hebrew, *Haaretz* 6.20.2003).
- c. A right wing Serbian National party has won **most** of the votes, but not enough to form a government (BBC World News, 12.29.2003).³¹

In multi-party systems, hardly ever does a party/candidate gain power without a coalition with a few other parties, which then together represent more than 50% of the voters. Since a candidate must receive at least 40% of the votes in order to avoid a second round of elections, and since s/he is expected to form a coalition, receiving 41.6% practically guarantees gaining more than 50% (by forming a coalition) in order to assume office (in a, and similarly, in c, presumably). Peres (in b) is the candidate who received the most votes, even though he got less than 50% of the votes, since there were a few, rather than two candidates. Hence, 49.2% counts as ‘most’.

³¹Presumably, the party does not have enough votes to form a government because it doesn’t have more than half of the votes.

Now, how can we explain this blatant violation of a meaning claimed to be coded by *most*? One option (proposed to me by Edit Doron, p.c.) is to assume that actually plurality, rather than majority is the coded meaning of *most*. This would mean that *most* means ‘the largest proper subset’, and the prevalent lower bound is actually pragmatically derived. Another option is to argue that *most* is lexically ambiguous between a majority meaning (lower- and upper- bounded, as specified above) and a plurality meaning (only upper bounded). Context then selects the appropriate meaning. A third possibility is that *most* is lower (and upper) bounded lexically, and it is the plurality interpretation which is pragmatically derived. While each of these solutions has some plausibility, they do not all account for the data equally well. I prefer the third option.

Assuming plurality to be the coded meaning of *most*, we can derive the lower bound by inference when context makes it clear that a binary partition is involved. Thus, while the largest subset may be smaller than 51% when there are more than two subsets, it must be at least 51% if there are only two subsets. It then seems quite straightforward to generate the lower bound contextually. All we need assume is that we (nowadays) tend to reduce partitions into one binary partition. This is why there seems to be some competition between the ‘largest subset’ meaning and the ‘above 50%’ meaning. Still, it seems that speakers accept the former only in some limited cases, but not in others (for details, see Ariel 2004:6.5). It is hard to imagine, for example, that Barclays Global Mutual Fund, which holds 3.535% of the shares of Ford Motor Company can be said to own *most of the shares*, even though it is the largest share holder of that company, and thus has the most shares.³² So it doesn’t look like being ‘the largest subset’ guarantees an appropriate use of *most*, which it should: *~Barclays Global owns most of the shares of Ford Motor Company* seems to express a false proposition (and not just an inappropriate utterance).

Because those cases where speakers accept as majorities percentages lower than 51% are rare, pragmatically restricted, and not consistent and stable enough, I favor the solution maintaining a lexical lower bound. The findings reported in Ariel (2003, 2004) demonstrate a categorical difference between the acceptance of 50% and 51% for *most* (but not between 49% and 50%, for example – see Figure 2), and attest to the hard-edged categorical nature of the lower bound. (Or, alternatively, to the automatic reduction into a binary partition contrasting some majority with some minority). This is why I tend to reject an ambiguity solution. If I am right, and the original plurality meaning has given rise to a lower (and upper-) bound majority meaning, we have to assume that where plurality (of subsets smaller than 51%) allows for the use of *most*, the linguistic meaning has been loosened up (I here adopt the Relevance-theoretic concept – see Carston 2002:Chapter 5). Note also that the journalist in (51b) explicitly mentions the fact that the quantity involved is less than 50%, which should not be noteworthy if *rov* simply refers to plurality. Interestingly, the same election result reported in (51a), when announced in the news (“Reshet Bet”, 6.4.2003), did not get labeled *rov*. However, the question of plurality vs. majority for *most* deserves further research.

6 Conclusions

In Ariel (2004) I have argued that what pragmatics cannot deliver (irrelevant “forced” ‘not all’ implicatures), semantics must (a lexical upper part). I have here argued that lexical semantics is indeed up to the challenge. I suggest that *most* means ‘a proper subset which is the largest subset, given any partitioning of the complement set (into one or more subsets)’. This definition entails

³²I thank Eric Berger, p.c. for finding such a case for me, which is valid for December 2003.

an upper (in addition to a lower) bounded lexical meaning, namely, that the quantity denoted by *most* is more than half and less than all. I have emphasized that the same **lexical** status should be attributed to the upper bound as to the lower bound for *most* (but not for *more than half*).³³ I believe that a circumbounded lexical meaning for *most* makes for a cognitively more reasonable concept to be coded by a single morpheme (Cf. the complex expressions for quantities lacking one bound: *more than half*, *at least half*, *at most half*).

Once we view the upper bound as lexically determined, the division of labor between lexical semantics and pragmatics is (partly) altered. I propose to shift some of the pragmatic burden to lexical semantics (a weaker version of the received view ‘not all’), and some of the semantic burden of *most* to pragmatic inferential processes (the confirmation of the compatibility with ‘all’). In addition, however, (pragmatic) implicatures are still responsible for classical pragmatic phenomena: The generation of ‘not all’ and ‘possibly all’ implicatures in a minority of cases where these are intended by the speaker, as well as for the understanding that the quantity denoted by *most* is noteworthy (Ariel 2004:see). Note that both ‘not all’ and ‘possibly all’ inferences may have two rather different cognitive statuses. The unmarked case is for the upper bound to apply just because the speaker denoted 51–99% (for short), so that amount alone is profiled. In the unmarked case ‘not all’ is true but irrelevant, because ‘all’ is irrelevant, and addressees don’t consider set members not denoted. But when a ‘not all’ implicature is generated (as I argued in Ariel 2004, this happens when there is a specific expectation for ‘all’), ‘not all’ is profiled. In this case ‘not all’ (as well as ‘all’) are relevant in the discourse, and addressees do consider set members not explicitly denoted. The same is true for ‘possibly all’ implicatures (when *most* is enriched to ‘at least most’). In such cases ‘possibly all’ is profiled. In the unmarked case, however, where no such implicature is generated, ‘possibly all’ is not profiled, and hence, a pragmatic inference is needed in order to view the speaker’s proposition as compatible with ‘possibly all’.

Note that on my view there is then a difference between implicating ‘possibly all’ or ‘not all’ and allowing for *most* to be compatible with ‘all’/‘not all’. As we have seen above (and see Ariel 2004 and Papafragou & Schwarz to appear), the latter pragmatic inferences are not guaranteed. Subjects are quite often reluctant to draw them, refusing to view a *most* utterance as compatible with ‘all’. This is not the case for when ‘possibly all’ is implicated. I doubt that there will be many addressees who will refuse to accept the ‘at least most, possibly all’ interpretation in ex. (4). But if implicated ‘possibly all’ and ‘not all’ are profiled in certain uses of *most*, what is the difference between these implicatures and explicit assertions of *possibly all* and *not all*? The difference is that the former can always be cancelled. The latter can only be taken back by a correction.

It is my hope that my work on *most* will encourage researchers to re-examine the division of labor between semantics and pragmatics for all the scalar expressions with respect to lower-bounded-only lexical meanings. I suggest that linguists apply the important distinction proposed by Koenig (1991) between truth-compatibility and lexical meaning, in order to avoid ‘surplus’ semantic analyses. Whereas lexical meanings must enable truth compatibility, they should not be viewed as the sole direct source for this compatibility. I believe that once we impose this distinction, we will offer more realistic lexical meanings, ones which account for actual speakers’ intended meanings. In the case of *most*, I have proposed that attending to this distinction yields a

³³Recall that an equally marginal percentage of responses were in violation of the lower and the upper bounds of *most* (see again Table 1).

'just that' semantic analysis, which should be preferred over an analysis where a default meaning is derived indirectly, via an implicature. I would not be surprised if other scalar quantifiers and predicates can also receive a 'just that' circumbounded lexical semantics.

Bibliography

- Anscombe, J.-C., Ducrot, O., 1983. *L'argumentation dans la langue*. Pierre Mardaga, Bruxelles.
- Ariel, M., 2002. Privileged interactional interpretations. *Journal of Pragmatics* **34**, 1003–1044.
- Ariel, M., 2003. Does most mean 'more than half'? *BLS* **29**, 17–30.
- Ariel, M., 2004. 'Most'. *Language* **80** (4), 658–706.
- Carston, R., 1990. Quantity maxims and generalized implicature. *UCL Working Papers in Linguistics* **2**, 1–31, revised version in 1995 *Lingua*, **96**, 213–244.
- Carston, R., 1998. Informativeness, relevance and scalar implicature. In: R. Carston, S. Uchida (eds.), *Relevance theory: Applications and interpretations*. John Benjamins, Amsterdam, 179–236.
- Carston, R., 1999. The semantics/pragmatics distinction: A view from relevance theory. In: K. Turner (ed.), *The semantics/pragmatics interface from different points of view*. Elsevier, Oxford, 85–125.
- Carston, R., 2002. *Thoughts and utterances*. Blackwell, Oxford.
- Chafe, W. L., 1994. *Discourse, consciousness, and time*. The University of Chicago Press, Chicago.
- Chierchia, G., 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In: A. Belletti (ed.), *Structures and beyond, the cartography of syntactic structures*. Vol. 3. Oxford University Press, Oxford, 39–103.
- Chierchia, G., McConnell-Ginet, S., 1990. *Meaning and grammar: An introduction to semantics*. MIT Press, Cambridge, Massachusetts.
- Geurts, B., 1998. Scalars. In: P. Ludewig, B. Geurts (eds.), *Lexikalische Semantik aus kognitiver Sicht*. Gunter Narr Verlag, Tübingen, 95–117.
- Giora, R., Balaban, N., Fein, O., Alkabets, I., 2004. Negation as positivity in disguise. In: H. L. Colston, A. Katz (eds.), *Figurative language comprehension: Social and cultural influences*. Erlbaum, Hillsdale, NJ, 233–258.
- Giora, R., Fein, O., Ganzi, J., Alkeslassy Levi, N., Sabah, H., 2005. On negation as mitigation: The case of irony. *Discourse Processes* **39**, 81–100.
- Grice, P. H., 1975. Logic and conversation. In: P. Cole, J. L. Morgan (eds.), *Syntax and semantics*. Vol. 3: Speech Acts. Academic Press, New-York, 41–58.
- Hamilton, W., 1860. *Lectures on logic*. Blackwood and Sons, Edinburgh.
- Heritage, J., 2002. Oh-prefaced responses to assessments. In: C. Ford, B. Fox, S. A. Thompson (eds.), *The language of turn and sequence*. Oxford University Press, Oxford, 196–224.
- Hirschberg, J. B., 1991. *A theory of scalar implicature*. Garland, New York.
- Horn, L. R., 1972. *On the semantic properties of logical operators in English*. Mimeo, Indiana University Linguistics Club, Bloomington, IN.
- Horn, L. R., 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In: D. Schiffrin (ed.), *Meaning, form, and use in context: Linguistic applications (GURT '84)*. Georgetown University Press, Washington, DC, 11–42.
- Horn, L. R., 1989. *A natural history of negation*. University of Chicago Press, Chicago.
- Horn, L. R., 1992. The said and the unsaid. In: Proceedings of the second conference on semantics and linguistic theory II. Dept. of Linguistics, Ohio State University, Columbus, 163–192.

- Horn, L. R., 1996. Presupposition and implicature. In: S. Lappin (ed.), *The handbook of contemporary semantic theory*. Blackwell, Oxford, 299–319.
- Horn, L. R., 2003. Assertoric inertia and NPI licensing. *CLS* **38** (Part 2: The panels), 55–82.
- Horn, L. R., 2004. Implicature. In: L. R. Horn, G. Ward (eds.), *The handbook of pragmatics*. Blackwell, Oxford, 3–28.
- Horn, L. R., 2005. The border wars: A neo-Gricean perspective. In: K. von Stechow, K. Turner (eds.), *Where semantics meets pragmatics*. Elsevier, Amsterdam (= this volume).
- Jespersen, O., 1949. *A modern English grammar on historical principles, part VII: Syntax*. Munksgaard, Copenhagen.
- Kadmon, N., 1987. *On unique and non-unique reference and asymmetric quantification*. Ph.D. dissertation, University of Massachusetts, Amherst, also published by Garland Press, 1992.
- Kadmon, N., 2001. *Formal pragmatics*. Blackwell, Oxford.
- Kay, P., 1989. *At least*. University of California at Berkeley, manuscript.
- Keenan, E. L., 1996. The semantics of determiners. In: S. Lappin (ed.), *The handbook of contemporary semantic theory*. Blackwell, Oxford, 41–63.
- Koenig, J.-P., 1991. Scalar predicates and negation: Punctual semantics and interval interpretations. *CLS* **27**, 140–155.
- Lasnik, R., 1999. Pragmatic halos. *Language* **75**, 522–551.
- Levinson, S. C., 2000. *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press, Cambridge, Massachusetts.
- Morris, C., 1994. *Bearing witness*. Little, Brown and Company, Boston.
- Newstead, S. E., 1995. Gricean implicature and syllogistic reasoning. *Journal of Memory and Language* **34**, 644–664.
- Noveck, I. A., 2001. When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition* **78**, 165–188.
- Noveck, I. A., Posada, A., 2003. Characterizing the time course of an implicature: An evoked potential study. *Brain and Language* **85**, 203–210.
- Papafragou, A., 2003. Scalar implicatures in language acquisition: Some evidence from modern Greek. In: *CLS 38*. University of Chicago Press, Chicago.
- Papafragou, A., Musolino, J., 2003. Scalar implicatures: Experiments at the semantics-pragmatics interface. *Cognition* **86**, 253–282.
- Papafragou, A., Schwarz, N., to appear. ‘Most’ wanted. *Language acquisition: Special issue*.
- Récánati, F., 1989. The pragmatics of what is said. *Mind and language* **4**, 295–329.
- Sadock, J. M., 1984. Whither radical pragmatics? In: D. Schiffrin (ed.), *Meaning, form, and use in context: Linguistic applications (GURT '84)*. Georgetown University Press, Washington, DC, 139–149.
- Sevi, A., 1994. *Exception phrases*. Seminar paper, Tel Aviv University, in Hebrew.
- Smith, C. L., 1980. Quantifiers and question answering in young children. *Journal of Experimental Child Psychology* **30**, 191–205.
- Sperber, D., Wilson, D., 1986/1995. *Relevance*. Blackwell, Oxford.

Appendix A

An alternative formalization of my proposal for a lexical meaning for *most* was kindly offered to me by Uriel Cohen (p.c.). According to his formulation, *most Q are P* means: *QP* is a proper subset of *Q* which is *P*, such that for every subset of *Q* which is disjoint from *QP*, *QP* is larger than it.³⁴ Compatibility with ‘all’ is inferable by the addressee when for some contextual reason awareness is drawn to the fact that ‘all’ may in fact be the case (see section 3). *QP* can then be chosen as the known subset of *Q* which is *P*, and the fact that the speaker did not explicitly rule out the possibility that ‘all’ is the case serves as a basis for an inference that the possibility that *Q-QP* (the complement of the reference set) is also *P* is not denied. This is why ‘possibly all’ can be consistent with a *most* utterance. Note, however, that compatibility with ‘all’ is distinct from a ‘possibly all’ implicature on my account (see also Ariel 2004). The former is not an interpretation intended by the speaker, the latter is.

Let us now see what happens when the speaker intends to convey more than the lexical meaning of *most* via implicatures (for a more detailed discussion see Ariel 2004). ‘Possibly all’ is generated as a Particularized Conversational Implicature, when the speaker intends to convey that she’s actually entertaining the possibility that ‘all’ (even though she only undertook upon herself to assert *P* of a subset of *Q* – see ex. (4)). In this case, then, there is no speaker commitment to the existence of a disjoint subset of *Q*. Next, if ‘all’ is contextually expected, a ‘not all’ implicature is derived from *QP* being a proper subset of *Q* combined with a comparison to the expected but not asserted ‘all’ (this is the default case under the received view, but is rather rare on my analysis – see Ariel 2004 for examples). In this case the speaker is actually ruling out the possibility that ‘all’, even though this implicature, like all implicatures, may be cancelled. Finally, the plurality ‘largest subset although less than half’ reading (see section 5) can be pragmatically derived when *Q-QP* (the complement of the reference set as a whole) is not very accessible (i.e., salient in context), while there are many accessible competing subsets of *Q* which are disjoint from *QP* and from each other. In other words, when the complement of the reference set denoted by *most* is not conceived of as one unified subset which would then count as the majority (a context-dependent phenomenon – see Ariel 2004), these disjoint subsets put together may be larger than *QP*. Crucially, however, this unification into one subset is not entertained, each subset disjoint from *QP* is considered separately, and each is smaller than *QP*.

³⁴Cohen himself, however, claims that there is no need for the “proper” addition to the formula, once he requires the “competing” subsets to be nonempty. It can then be derived referentially, by disallowing *every* to quantify over empty domains. The inappropriateness of ‘all’ can be accounted for by the proposition being vacuously true, which is pragmatically dispreferred. His analysis can be seen in <http://www.tau.ac.il/~cohenfr>.