

Numerical expressions, implicatures and imagined prior context

Chris Cummins

Department of Linguistics and English Language

University of Edinburgh

c.r.cummins@gmail.com

Abstract

Pragmatic interpretations are, by definition, influenced by contextual factors. Research in experimental semantics and pragmatics has suggested that participants, when presented with fragments of discourse, draw inferences about the nature of the prior context and use these inferences to shape their interpretation of the target materials. This has both methodological and theoretical implications. Focusing on the domain of numerical expressions, I present an experiment that aims to elucidate the effect of participants imagining a particular prior context (specifically, one in which a given numeral is salient). I show that this expectation influences pragmatic interpretation in a classically predictable way. I further argue that the effect of ‘imagined prior context’ might be responsible for a sizeable portion of the unexpected variability exhibited between participants in typical pragmatic experiments.

1 Introduction

A substantial body of research in experimental semantics and pragmatics has addressed the generation of so-called scalar implicatures (SIs). SIs constitute a special case of the more general quantity implicature, in which – following the analysis of Grice (1989) – hearers use the speaker’s utterance to draw inferences about the falsity of logically stronger alternatives that could have been uttered instead. SIs specifically rely on the existence of informational scales, comprising terms which belong to the same semantic field but differ in informational strength.

The canonical example of scalar implicature, both historically and in the current experimental literature, involves the scale <some, all>. Taking “some” to possess purely existential semantic meaning, “all” entails “some”, and in that sense is informationally stronger (across a wide range of possible contexts of use). Consequently, the hearer of (1) is argued to be able to recover the implicature (2), as first observed by Mill (1865).

(1) I saw some of your children today.

(2) The speaker saw some but not all of the addressee’s children today.

The availability of such an implicature relies upon a number of auxiliary assumptions, including that the speaker is knowledgeable about the stronger proposition (as already pointed out by Mill) and potentially that the stronger proposition is relevant to the discourse purpose (see for example Breheny et al. 2006). However, those assumptions being met, implicatures should be recoverable by any competent user of language. Indeed, on a Gricean analysis, they are an aspect of intentional communication: the speaker of (1) explicitly intends to convey the meaning “some but not all”, and the work of the hearer is merely to recover this intention. In that sense, the ability to recover implicatures is a necessary part of a language user’s communicative competence (at least if we accept the general characterization of linguistic communication as ‘intentional’).

From this point of view, it is unsurprising that developmental research has documented that young children appear to lack facility with implicatures (Papafragou & Musolino 2003, Guasti et al. 2005, and many others). However, it is profoundly surprising that numerous adult studies have documented acceptance rates for the implicature “some” +> “not all” that are far from maximal (Noveck 2001, Bott & Noveck 2004, Guasti et al. 2005, etc.)

This cannot readily be attributed to deficiencies in the specific scale being tested, <some, all>. Of course, this scale may indeed be deficient in some respect, but comparative research suggests that it is nevertheless among the strongest and most reliable of the posited implicatural scales (van Tiel et al. in prep.) Hence, if the <some, all> scale lacks explanatory value, we might argue that the same is true of scalar implicature in general.

A less radical alternative account for the variability in performance, both between and within tasks, is that it is driven by contextual factors. Depending on the precise nature of the task, an underinformative choice of expression – such as

saying “some” when in fact “all” is the case – might be acceptable to a greater or lesser degree. For instance, we might expect that underinformative “some” would be less acceptable if the task is understood to involve giving the best possible description, but more acceptable if the task merely involves making any true statement. The nature of the judgment that participants are obliged to make could also exert an influence here, as for instance in Katsos and Bishop’s (2011) study. They demonstrate that children aged 5 reliably accept (and adults reliably reject) descriptions with “some” given to situations with “all”, when the response condition is effectively binary (yes/no). However, when the response condition is ternary (in effect, good, bad or medium), children and adults alike reliably assign the intermediate rating to underinformative descriptions with “some”. This suggests, as Katsos and Bishop argue, that the children’s behaviour in the binary condition does not reflect their lack of awareness of the shortcomings of the tested utterances. Rather, it seems to reflect an unwillingness on the children’s part to reject utterances on this basis, an unwillingness that adults do not share.

Can we invoke a contextual explanation to deal with within-task variability, though? In such cases, the presented context is the same for all participants, yet the observed behaviour varies. The only possible contextual explanation for this is that participants – in addition to taking into account the provided context – are imagining more elaborate and detailed prior contexts for the utterances, and that these contexts differ between participants, for instance in the level of accuracy or informativeness that they require the following utterance to exhibit.

The idea that participants in experiments of this kind might conjure up richer contexts for interpretation is not a new one – Breheny et al. (2006), for instance, explicitly note this possibility. However, it appears that relatively little attention has been paid to documenting directly whether this phenomenon exists, and if so, whether or not it is widespread. This omission is surprising given the potential methodological importance of such work for experimental semantics and pragmatics. As a research area, experimental pragmatics grapples directly with this issue, in that the object of study is the meanings of real-life utterances produced in particular contexts, but the experimental research that addresses this question relies heavily on artificially constructed materials which are necessarily often

presented in relatively impoverished contexts. In experiments, it is more typical to present a single conversational turn or a question-answer pair than a full dialogue, and it is hard to exclude the possibility that participants may make assumptions about the higher-order discourse purpose or the content of previous turns to which they were not privy.¹

Indeed, even our theoretical intuitions about pragmatic meanings may be informed by speculation about the likely context of utterance, even when this is not treated in a systematic fashion by theory. Even the uncontroversial intuition that “some” can convey “not all” relies on the assumption that the stronger proposition “all” might have been relevant, given the prior discourse context, in circumstances in which “some” can be uttered, an assumption that in turn relies on a notion of relevance that is somewhat elusive. For less frequently occurring forms, such as those discussed in the following section, the problem may be more severe, as the form may effectively carry more information about its own likely context of utterance than is generally acknowledged.

In this paper, I make a preliminary attempt at addressing the issue of ‘imagined prior context’ experimentally. In doing so, I focus on pragmatic enrichments within the numerical domain, a decision that I attempt to motivate in the following section.

2 Implicatures from numerical expressions

The domain of numerical expressions appears to be a fertile one for pragmatic enrichment. A popular analysis of numeral meaning holds that numbers are lower-bounded on their semantics and acquire exact meanings pragmatically through implicature (although see Breheny 2008 for a critical discussion of this proposal). More recently, Cummins, Sauerland and Solt (2012) demonstrate the availability of pragmatic enrichments, apparently due to quantity implicature, from expressions of the form “more than n ”.

¹ An anonymous reviewer raised the general and very important question of what artificial experiments of this kind can tell us about natural communication. I have no space here to offer a manifesto for experimental pragmatics, as practised at the sentence level. However, I would argue that both the process of enriching weak scalar meanings and the process of inferring non-shared prior context are highly likely to be relevant to natural communication. Nevertheless, my immediate concern here is just to try to disentangle those two processes in laboratory tasks.

They also argue that these enrichments are conditioned by numeral salience.

To take a specific example, Cummins et al. (2012) show experimentally that quantifying sentences such as (3) are considered to convey additional meanings to the effect that, for instance, (4) or (5).

- (3) I have more than 60 CDs.
- (4) I do not have more than 80 CDs.
- (5) I do not have more than 100 CDs.

The available implicatures are argued to depend upon the salience of the numeral concerned. That is, Cummins et al.'s account explains the absence of an implicature to the effect that (6) is false, given the utterance (3), by arguing that (6) is independently disfavoured on the basis of using a non-salient number. Hence, the speaker's decision to utter (3) rather than (6) can be explained just as a preference for using the number 60 rather than 61, and consequently there is no need for the hearer to postulate that the speaker is unable to commit to the truth of the assertion (6). For this reason, the implicature not-(6) is predicted to be unavailable, as is borne out experimentally.

- (6) I have more than 61 CDs.

Whether or not this particular account is along the right lines, Cummins et al.'s data seems strongly to suggest that implicatures are available in principle from utterances containing "more than n " for numeral n . Moreover, for certain values of n , a wide range of different implicatures appear to be available, depending on the preferences of the individual participant. A given instance of "more than 100" can be construed as conveying "not more than 110", "not more than 125", "not more than 150" or "not more than 200". Hence, just like the some/all case, there is considerable variation between participants as to whether specific pragmatic enrichments are endorsed. Indeed, the picture is more colourful in the numerical case, inasmuch as a greater number of distinct candidate implicatures (or sets of implicatures) are endorsed by different participants, but again the reasons for this are not clearly understood. Moreover, as noted by Fox and Hackl (2006), such implicatures are not observed in the cases of small cardinal quantities ("more than two people" does not implicate "not more than three people"), which is another fact requiring explanation.

For numerical expressions, as opposed to other expressions of quantity, it also seems more feasible to be able to ask participants direct questions about the choice of expression. Given an utter-

ance such as (3), the question "Do you think that the specific number 60 was important for some reason?" seems perfectly reasonable and is not a leading question. By contrast, given an utterance such as (1), the question "Do you think that the specific quantity 'some' was important for some reason?" seems less natural.

For all these reasons, I would argue that the domain of numerical expressions is a particularly convenient testbed for the hypothesis sketched out in the introduction: namely that the variability between participants in their generation of implicatures is partly explicable in terms of the different prior contexts that they imagine. The experiment in the following section sets out to investigate this claim.

3 Experiment: implicatures and inferences about prior context

In this experiment, participants read sentences containing numerically-quantified expressions, and were asked a set of questions about each sentence. The aim was to examine simultaneously whether the kind of implicature predicted by Cummins et al. (2012) was available, whether the reader inferred that the specific number was being used for a particular reason, and whether (as predicted by, for instance, a traditional Gricean pragmatic account) these two forms of inference were inversely correlated in strength.

3.1 Materials

12 sentences containing numerically-quantified expressions were sampled from the BNC (BNC, 2007). These comprised one instance each of "more than 60", "more than 70", "more than 80", "more than 90", "at least 60", "at least 70", "at least 80", "at least 90", "more than one", "more than two", "more than three", and "more than four". The usage of each expression was cardinal and related to the number in question: instances such as "more than 50 per cent", "more than 60 million", and "more than 70 metres" were excluded from consideration. Bearing in mind Cummins et al.'s (2012) findings about the presence of prior context, sentences were also excluded from consideration if the preceding sentence contained a numeral (or if there was no preceding sentence, i.e. the sentence in question was the beginning of a text). However, the preceding sentences were in any case not presented to participants in this study.

Instances of “more than/at least n ” for non-round n are rare in the BNC and no appropriate examples of cardinal usage, respecting the above criteria, could be located. For this reason, non-round conditions were created by replacing the above numbers with non-round numbers of the same order of magnitude: 60 with 58, 70 with 77, 80 with 86, and 90 with 93.²

Two lists were created, each comprising 12 items in pseudorandom order. The four small-number “more than” sentences were presented on both lists. For the remaining items, the design balanced between round (original) and non-round (replacement) numbers. Thus, version 1 contained sentences with “more than 60”, “more than 77”, “more than 86” and “more than 90”, whereas version 2 contained those same sentences with “more than 58”, “more than 70”, “more than 80” and “more than 93”. For “at least”, the reverse was true: version 1 contained “at least 58/70/80/93” and version 2 contained “at least 60/77/86/90”. In this way, each participant saw each sentence and each number only once. The sentences used are shown in Appendix A.

For each item, participants were asked to judge four statements on a five-point Likert scale rated from “very unlikely” (1) to “very likely” (5). The first statement concerned the availability of a specific implicature predicted by Cummins et al. (2012); for instance, where the text identified the existence of “more than 70 volumes”, statement (i) was “In the speaker’s opinion, the actual number of volumes is less than 80”. Statement (ii) was “The speaker said [more than 70] because that was the most informative statement possible”. Statement (iii) was “The speaker said [more than 70] because that was a convenient approximation”. Statement (iv) was “The speaker said [more than 70] because the specific number [70] was important for some reason”.

3.2 Participants

Participants were recruited via Amazon Mechanical Turk. The conditions were fielded on separate days in April 2014. 17 participants completed version 1 of the experiment and 14 participants completed version 2.

² An anonymous reviewer observes that the construction of materials in this way could be seen as an advantage, in that it reduces the amount of irrelevant variance. However, for the present purposes, I consider this a potential disadvantage, as I must then assume without proof that the resulting materials are in fact pragmatically felicitous.

3.3 Results

As no major differences were observed between the results from the two conditions, they are pooled and considered together in what follows. Table 1 presents the mean ratings (and SDs) for each of the test conditions.

	(i)	(ii)	(iii)	(iv)
More than				
Round	3.46 (1.30)	3.44 (1.15)	4.08 (1.06)	2.98 (1.09)
Non-round	3.63 (1.12)	3.68 (1.04)	3.29 (1.23)	3.11 (1.27)
Small	2.02 (1.27)	3.43 (1.13)	3.29 (1.20)	3.58 (1.24)
At least				
Round	3.37 (1.41)	3.67 (1.04)	3.90 (0.94)	3.10 (1.16)
Non-round	3.27 (1.38)	3.87 (1.09)	3.21 (1.33)	3.27 (1.26)

Table 1: Mean ratings (and SDs) for each quantifier and number condition

Considering the mean responses for each tested item within each category (i.e. the means by-sentence), the ratings for (i) and (iv) are strongly negatively correlated (Pearson’s $r = -0.67$). These mean ratings are tabulated in full in Appendix B. Planned comparisons via t-tests indicate that the ratings in the “more than” condition with respect to statement (i) are lower for small numbers than for either round or non-round numbers, and with respect to statement (iv) are higher for small numbers than for either round or non-round numbers (all $p < 0.01$).

3.4 Discussion

The existence of a strong negative correlation between judgments of statements (i) and (iv) seems to suggest that, where participants infer that specific numerals are being used for a particular reason, they are disinclined to infer the otherwise-predicted pragmatic enrichment. This appears to concur with the predictions of Cummins et al. (2012). Recall that the availability of an enrichment of the kind canvassed in (i) requires that a stronger alternative assertion was available to the speaker, and that this alternative was not selected purely on the grounds of its falsity. By contrast, where a specific numeral is

chosen because it is somehow intrinsically special (as evidenced by a high rating for statement (iv)), the informationally weaker assertion may be preferable to informationally stronger alternatives, on the basis that these stronger alternatives would fail to use the “special” number. Consequently, the speaker’s decision to use the informationally weaker assertion should not convey anything about the truth-value of the informationally stronger alternative in this particular case.

Delving into the specific conditions, the results suggest that participants are strongly disinclined to endorse the candidate implicatures arising from the small number conditions “more than two/three/four/five” (respectively, “not more than three/four/five/six”). This is unsurprising – these implicatures have been widely assumed to be unavailable (see for example Fox and Hackl 2006), at least in cardinal contexts. More strikingly, these expressions give rise to clear judgments that the numbers in question are likely to be contextually salient (as shown by their high ratings on statement (iv)), even in the absence of any explicit contextual support for this claim.

The unavailability of these implicatures could be attributed to several distinct causes. One possibility (explored by Fox and Hackl 2006) is that expressions of the form “more than n ” systematically fail to give rise to implicatures: however, this appears to over-predict, in the light of Cummins et al.’s data. Another possibility is that the implicatures are blocked as a consequence of their communicative oddness: if “more than two” implicated “not more than three”, these premises would together entail “exactly three”, which could be much more easily communicated in other words. This would also account for the intuition that “more than two” gives rise to implicatures in measurement contexts, with “more than two metres” implicating “not more than three metres”. However, the results of this experiment could be taken to support a third explanation, namely that the systematic lack of implicatures from expressions such as “more than two” stems from the fact that these expressions trigger strong expectations that the specific numeral used was used for a particular reason. A rational hearer who held such an expectation should be unwilling to draw quantity implicatures. For instance, suppose that the hearer assumes “more than two” is being used because “two” is an especially salient number. It follows that the more informative “more than three” might not be a better alternative, even if it is true, on the basis

that it fails to use this salient number “two”. The hearer should conclude that the use of “more than two” rather than “more than three” does not necessarily signal the speaker’s unwillingness to commit to the truth of that latter, stronger proposition.

Of course, this explanation is only tenable if sentences involving “more than two” in cardinal contexts are restricted in their distribution. They would be predicted to be admissible in situations in which the number “two” is salient, or can be presumed to be salient: in such situations, the implicature “not more than three” would be blocked for the reason discussed above. “More than two” would also be predicted to be admissible in situations in which the speaker is not knowledgeable about the truth of stronger propositions, in which case the implicature would fail to arise for standard reasons (this epistemic assumption being essential for implicature on the traditional account). However, “more than two” would be predicted not to be admissible in situations in which the speaker is knowledgeable about the precise value and in which the number “two” is not especially salient. Examples discussed in the literature such as (7), in which the speaker turns out to be knowledgeable about the precise value, appear strongly to invite the inference that having “two children” constitutes a threshold of some kind (e.g. for entitlement for benefits). However, the question remains open as to whether all examples of “more than two” in cardinal quantificational contexts actually have this property.

(7) John has more than two children; in fact, he has five.

In the case of large round numbers, participants are inclined to draw the pragmatic enrichment, endorsing statement (i). This replicates the findings of Cummins et al. (2012). Moreover, participants strongly endorsed statement (iii) in this case (the rating exceeding that for both other conditions; t-tests, $p < 0.01$). This suggests that these utterances are regarded as convenient approximations rather than attempts to use specific numbers; hence, implicatures should be available. This expectation seems to be borne out.

Large non-round numbers behave similarly to large round numbers in this experiment, but were numerically rated higher with respect to both statement (iv) and statement (i). They scored somewhat lower on (iii), perhaps indicating that they are not as ‘convenient’ an approximation as round numbers; and slightly higher on (ii), suggesting that they can be perceived as optimally

informative. This fits with the assumption that the use of non-round numbers permits greater precision but is associated with additional cognitive costs. It is tempting to hypothesize that the large non-round numbers constitute an intermediate case between round and small numbers in this experiment, and that the speaker who uses such a number is presumed both to be deliberately using a specific number and to be attempting to convey an implicature. This would be conceivable if the hearer presumes that the speaker might prefer to use some specific number, but may not be willing to sacrifice a great deal of informativeness in order to do so: for example, even if 83 is a salient number, a speaker might use “more than 100” in preference to “more than 83” if they know the informationally stronger statement to be true. However, more work is required both in order to determine whether speakers actually exhibit this kind of preference, and – independently of that – whether hearers perceive that speakers are going to exhibit this kind of preference, and can modulate their interpretations of quantity expressions accordingly.

4 Conclusion

The experiment presented in this paper represents a preliminary attempt to explore the idea that numerically-quantified expressions might signal information about the prior context against which they should be interpreted, even when this prior context is not provided. The results of the experiment do appear to suggest that this is the case: participants spontaneously infer that specific numbers (of particular kinds) are contextually salient, purely on the basis of their usage. The implicatures recovered by participants appear to be modulated by this perception of contextual salience, although it is not possible to infer the existence of a causal relationship on the basis of this experiment.

Based on these findings, it is tempting to posit that at least some of the variability between participants, documented in experiments on quantity implicature, might be attributed to differences in the way in which they infer details of the context of utterance. The domain of number represents a convenient testbed for this approach, but in principle the hypothesis makes predictions about a much wider range of situations. Future work will aim both to broaden and deepen the experimental exploration of this area.

References

- BNC. 2007. *The British National Corpus*, version 3 (BNC XML Edition). 2007. Distributed by Oxford University Computing Services on behalf of the BNC Consortium. <http://www.natcorp.ox.ac.uk/>.
- Lewis Bott and Ira A. Noveck. 2004. Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of Memory and Language*, 51(3): 437-457.
- Richard Breheny, Napoleon Katsos and John N. Williams (2006). Are generalised scalar implicatures generated by default? An on-line investigation into the role of context in generating pragmatic inferences. *Cognition*, 100(3): 434-463.
- Richard Breheny. 2008. A new look at the semantics and pragmatics of Numerically Quantified Noun Phrases. *Journal of Semantics*, 25(2): 93-139.
- Chris Cummins, Uli Sauerland and Stephanie Solt. 2012. Granularity and scalar implicature in numerical expressions. *Linguistics and Philosophy*, 35: 135-169.
- Danny Fox and Martin Hackl. 2006. The universal density of measurement. *Linguistics and Philosophy*, 29: 537-586.
- H. Paul Grice. 1989. *Studies in the Way of Words*. Harvard University Press, Cambridge, MA.
- Maria Teresa Guasti, Gennaro Chierchia, Stephen Crain, Francesca Foppolo, Andrea Gualmini and Luisa Meroni. 2005. Why children and adults sometimes (but not always) compute implicatures. *Language and Cognitive Processes*, 20(5): 667-696.
- Napoleon Katsos and Dorothy V. M. Bishop. 2011. Pragmatic tolerance: implications for the acquisition of informativeness and implicature. *Cognition*, 120: 67-81.
- John Stuart Mill. 1865. *An examination of Sir William Hamilton's philosophy and of the principal philosophical questions discussed in his writings* (2nd ed.). Longmans, Green and Co., London.
- Ira A. Noveck. 2001. When children are more logical than adults: Investigations of scalar implicature. *Cognition*, 78: 165-188.
- Anna Papafragou and Julien Musolino. 2003. Scalar implicatures: experiments at the semantics/pragmatics interface. *Cognition*, 86: 253-282.
- Bob van Tiel, Emiel van Miltenburg, Natalia Zevakhina and Bart Geurts. In prep. Scalar diversity.

Appendix A. Materials, including variant numbers used

Materials used in this experiment have been extracted from the British National Corpus, distributed by Oxford University Computing Services on behalf of the BNC Consortium. All rights in the texts cited are reserved.

1. These are supplemented by more than 60/58 massive volumes of press-cuttings. (BNK 829)
2. We just hit at the right moment and from that week onwards, at least 93/90 people turned up. (AB5 566)
3. You may also have noticed that there are more than four grey shades used. (CGU 967)
4. They have lured or chased more than 77/70 species of vertebrates around racetracks in the Kenyan desert, up treadmills at the field station, and over runways of force plates in Milan, all in the interests of learning, as it were, how many kilometres each model gets per litre. (B75 1009)
5. In December 1984 at least 80/86 Jehovah's Witnesses were arrested in Limbé, southwest Cameroon, after holding an unauthorized religious meeting. (A03 628)
6. Violence was believed to be declining; the last war involving more than two great powers had been fought in the Crimea, far away, and the assumptions which governed fighting were more humane than ever before. (CM6 1021)
7. Plant experts at the meeting of the Convention on International Trade in Endangered Species (CITES) have agreed that more than 86/80 species of 'slipper' orchids — the genus *paphiopedilum* from Asia and the genus *thragmipedium* from South America — should be listed on the CITES Appendix I, which bans all commercial trade. (A59 421)
8. At least 70/77 alternatives have been submitted, with that of "Polish Socialist Labour Party" the front-runner. (A7V 300)
9. In the next example the character's thought spans more than one sentence. (EF8 1488)
10. Iranian-born Khoei, a scholar who had written more than 90/93 books on theology, was known for his adamantly apolitical stance. (HLN 2053)
11. On the basis of earlier work relying on measuring footprints, it had been estimated that

there must be at least 58/60 rhinos in the park. (J3K 92)

12. We only have to look at Tintswalo Hospital (Gazankulu) and more than three surrounding villages that fall under the jurisdiction of Lebowa Authority for evidence of this inaccessibility. (FBH 1174)

Appendix B. Mean ratings by-sentence

Tables 2 and 3 present the mean ratings for each sentence in versions 1 and 2 of the experiment. Sentences are numbered as in Appendix A; where applicable, the first-given number in Appendix A was used in version 1 of the experiment, and the second-given number was used in version 2 of the experiment.

Sentence	(i)	(ii)	(iii)	(iv)
1	2.94	3.29	4.06	3.06
2	3.59	3.88	2.94	3.53
3	1.82	3.65	3.82	3.71
4	3.18	3.88	3.18	3.47
5	3.59	3.59	3.76	3.29
6	1.88	3.53	3.59	3.82
7	3.76	3.71	3.35	3.35
8	3.24	3.59	3.94	3.29
9	2.06	3.29	3.18	3.29
10	3.47	3.65	4.29	3.00
11	2.59	4.00	3.29	3.47
12	2.06	3.71	3.29	3.53

Table 2: Mean results by-sentence in version 1 of the experiment

Sentence	(i)	(ii)	(iii)	(iv)
1	4.07	3.57	3.29	3.14
2	3.14	3.64	4.00	3.07
3	1.86	3.36	3.43	3.79
4	3.46	3.50	3.93	3.07
5	3.79	3.93	3.07	2.93
6	2.43	3.21	2.79	3.64
7	4.07	3.29	4.00	2.79
8	3.21	3.64	3.57	3.07
9	2.14	3.14	2.93	3.64
10	3.57	3.50	3.36	2.36
11	3.50	3.92	3.93	2.64
12	2.00	3.43	3.14	3.21

Table 3: Mean results by-sentence in version 2 of the experiment