

TOWARDS A MODULAR APPROACH TO ANAPHOR RESOLUTION

Arnout W. Koornneef, Frank Wijnen & Eric Reuland
Utrecht Institute of Linguistics OTS
Utrecht University
e-mail: [arnout.koornneef \[at\] let.uu.nl](mailto:arnout.koornneef@let.uu.nl)

Abstract

Recent accounts of anaphor resolution propose a two-route architecture for the interpretation of pronominals (e.g., Grodzinsky & Reinhart, 1993; Reuland, 2001). A pronoun can either be resolved by a grammatical operation in *logical syntax* (i.e., variable binding) or through value assignment in *discourse* (i.e., co-reference). Reuland (2001) proposes that an interpretation through variable binding requires less processing resources and is therefore preferred over a co-reference interpretation. Rule I compares variable binding with co-reference interpretations to decide whether a co-reference dependency is allowed (Grodzinsky & Reinhart, 1993; Reinhart 2000). This rule prevents discourse processes from by-passing logical syntax where the latter rules out an interpretation as ungrammatical. The question is whether Rule I always compares variable binding with co-reference interpretations, or is only executed if a dependency is initially ruled out by logical syntax. In an eye-tracking experiment we manipulated the interpretation of ambiguous and unambiguous pronouns in an attempt to explore how variable binding, co-reference and Rule I influence the way readers resolve pronouns. The results show that if a pronoun was ambiguous between a variable binding and co-reference antecedent, the variable binding antecedent was initially preferred even if discourse information clearly favoured the co-reference antecedent. Therefore, we argue that logical syntactic processes function independently from discourse processes, indicating that it may be warranted to consider logical syntax and discourse as distinct modules of the language system. Furthermore, the language system does not seem to apply Rule I in cases where both variable binding and co-reference lead to the same grammatical interpretation suggesting that Rule I is only relevant in cases that are potentially ungrammatical.

Background

Experimental studies have identified a range of factors that influence the resolution process for pronouns. For instance, subjecthood or first mention (e.g., McDonald & McWhinney, 1995), gender information, (e.g., Arnold, Eisenband, Brown-Schmidt & Trueswell, 2000), recency or distance (e.g., Clark & Sengul, 1979), and implicit causality information of interpersonal verbs (e.g., Caramazza, Grober, Garvey & Yates, 1977; Koornneef & Van Berkum, 2006) all seem to affect the search for a referent. Within psycholinguistics the current debate mainly focuses on the question *when* during comprehension these factors influence the pronoun resolution process. For instance, some argue that gender information becomes available and is put to use before other factors can have an affect (e.g., Ehrlich, 1980), while others claim that gender is only used during special, strategic or later processing (e.g., Greene, McKoon & Ratcliff, 1992). Still others propose fully dynamic

accounts where multiple sources of information are used in parallel to guide the resolution process (e.g., Arnold et al.). However, these different views have in common that they all implicitly assume that the pronoun is connected to a referent by accessing the discourse representation of the preceding written or spoken text. More specifically, they presume a one-route architecture to resolve a pronoun.

On the other hand, linguistic accounts of pronoun resolution (e.g., Reinhart, 1983) argue that the language system has two ways by which a pronoun is connected to an antecedent; (i) the pronoun behaves as a variable and is *bound* by its antecedent, or (ii) the pronoun receives a value from the discourse storage through co-reference. Consequently, sentence (1) has actually two possible representations, one in which ‘*he*’ is bound by ‘*the clown*’ (ex. 1a) and one in which the value of ‘*he*’ can be freely chosen from the discourse. If in the latter ‘*the clown*’ is picked as the proper antecedent for the pronoun the two derivations have the same interpretation (ex. 1b).

- (1) The clown_{*i*} thinks that he_{*i*} is funny
 (a) The clown λx (x thinks x is funny)
 (b) The clown λx (x thinks a is funny) and $a = x$

Variable binding is only possible if the antecedent c-commands the pronoun. Informally, a c-commanding antecedent can be characterized as ‘being higher in the syntactic tree’ of a sentence than the pronominal.¹ As a result, variable binding is licensed in example (1) but not in example (2). In the former ‘*the clown*’ is higher in the syntactic tree than ‘*he*’, but in the latter there is no c-command relation because the antecedent and pronoun appear in different sentences. Hence, in structures like (2) only a co-referential dependency can be constructed. On the other hand, if the antecedent contains a quantifier such as ‘*every*’, co-reference is ruled out (see ex. 3) and variable binding is in fact the only option for establishing an anaphoric dependency (e.g., ex. 4).

- (2) The clown_{*i*} knows it for sure. He_{*i*} is funny.
 (3) *Every clown_{*i*} knows it for sure. He_{*i*} is funny.
 (4) Every clown_{*i*} thinks that he_{*i*} is funny

The two-route architecture of the system comes with one important implication illustrated in (5). Logical syntactic constraints (i.e., Principle B) rule out a variable binding relation between the pronoun and the antecedent, as the pronoun is structurally too close to the antecedent (i.e., not ‘free’ in its binding domain). However, now we still have to explain why the availability of a co-referential interpretation does not systematically bypass the effect of Principle B.

- (5) *The clown_{*i*} hates him_{*i*}.

¹ A widely accepted formal definition of c-command is as follows: phrase α c-commands phrase β if and only if phrase α does not contain phrase β and the first branching node dominating phrase α also dominates phrase β (see Reinhart, 1983, for discussion).

A solution is provided by Rule I.² The original idea was that Rule I always compares the variable binding interpretation to the co-reference interpretation and opts for the former unless the two potential dependencies yield different interpretations (Grodzinsky & Reinhart, 1993). By applying this rule to (5), both variable binding and co-reference are ruled out; variable binding is blocked through Principle B and co-reference is therefore blocked as well because the two interpretations would be identical (i.e., in both cases ‘*he = the clown*’). However, on this construal Rule I does not only bear on ill-formed dependencies like (5), but on all dependencies that have an antecedent in a potential binding position, and a co-reference solution that points to that same antecedent. Thus, in grammatical example (1), where both variable binding and co-reference are possible in principle, the costly application of Rule I is predicted as well. Recently, however, Reinhart (2000) has proposed that Rule I only prevents co-reference from bypassing variable binding when the grammar rules out a variable binding interpretation and, hence, is only executed in potential ungrammatical cases like (4) and not in cases like (1).

According to Reuland (2001), variable binding and co-reference result from distinct processes in two functionally independent modules. Furthermore, it is proposed that a general economy principle governs the division of labour between the modules, based on the number of cross modular steps. That is, variable binding is thought to be more economic than co-reference because it requires fewer steps. By incorporating the economy principle, the theory assumes a processing hierarchy in which variable binding has precedence over co-reference. Hence, if a pronoun is ambiguous between a variable binding and co-reference antecedent the pronoun is predicted to be (initially) interpreted as referring to the former and not the latter. Given the logic of the theory this preference may stem from two sources; (i) differences in processing cost intrinsic to the processes themselves, or (ii) a time course effect – the variable binding option becomes available before the co-reference option and, hence, is initially taken. A similar proposal was made by Frazier and Clifton (2000) who hypothesize that ‘a bound-variable interpretation is preferred because the perceiver need only consult the Logical Form representation (not the discourse representation) in order to identify the bound-variable analysis of the sentence’. In a series of reading time and questionnaire experiments they examined this hypothesis and reported mixed results. It seemed that consistent results could only be found in VP-ellipsis contexts and did not generalize to other (quantificational) contexts, which led them to reject their ‘LF-only’ hypothesis.

In this study, we addressed the hypothesis that the parser prefers a variable binding over a co-reference dependency using structures very different from VP-ellipses. In addition, we studied whether the preference as implied by Reuland’s model might be attributed to an intrinsic difference in the associated processing costs or is rather a time

² A simplified version of original Rule I is: a pronoun cannot co-refer with a particular antecedent if the co-reference interpretation is indistinguishable from what would be obtained if the antecedent binds the pronoun.

A simplified version of revised Rule I is: a pronoun cannot co-refer with a particular antecedent if (i) the antecedent is in a potential binding position, yet cannot bind the pronoun (due to Principle B), and (ii) the co-reference interpretation is indistinguishable from what would be obtained if the antecedent binds the pronoun.

It is important to note that the original and revised Rule I are formulated in a way that correctly allows co-reference to by-pass Principle B in some specific situations. For instance, the dependency ‘*him = Bill*’ is allowed in the following example: “It is clear what Bill and Mary have in common, she adores him, and Bill_i adores him_i, too”.

course effect. Finally, we examined whether the revised version of Rule I should be preferred over the original version, or alternatively, that we should maintain the latter.

An eye-tracking experiment

In an eye-tracking experiment 36 Dutch participants read a series of short stories containing 36 experimental items (latin square design). In the experiment we manipulated the resolution of ambiguous and unambiguous pronouns to examine the issues discussed above.

The hypothesis that logical syntactic operations are preferred over discourse operations predicts that if a pronoun is ambiguous between a variable binding and co-reference antecedent the language system initially constructs a variable binding dependency, even if the context strongly supports a co-reference dependency. To test this prediction we presented stories like (6) and (7). In both stories the critical sentence contains the ambiguous pronoun 'he' with two potential antecedents (i.e., 'every worker' and 'Paul'). In story (6) the context preceding the pronoun supports a reading in which *he* refers to the variable binding antecedent (i.e., '*he = every worker*'). In story (7), on the other hand, the preceding context strongly supports a co-reference reading (i.e., '*he = Paul*').³ Reuland's model predicts that in story (7) initially the 'wrong' (variable binding) antecedent is chosen, because this dependency is initially preferred over the co-reference dependency. As a result, in (7) readers have to re-analyze their initial interpretation to construct the more suitable co-reference dependency.

(6) *Biased towards variable binding antecedent (VB-bias)*

A working day in the factory is always very tough. Especially today a lot of workers, among them the very old man Paul, could barely cope. **Every worker**_i who just like Paul was running out of energy, thought it was very nice that **he**_i could go home early this afternoon. After a hot shower things would probably look better.

(7) *Biased towards coreference antecedent (CR-bias)*

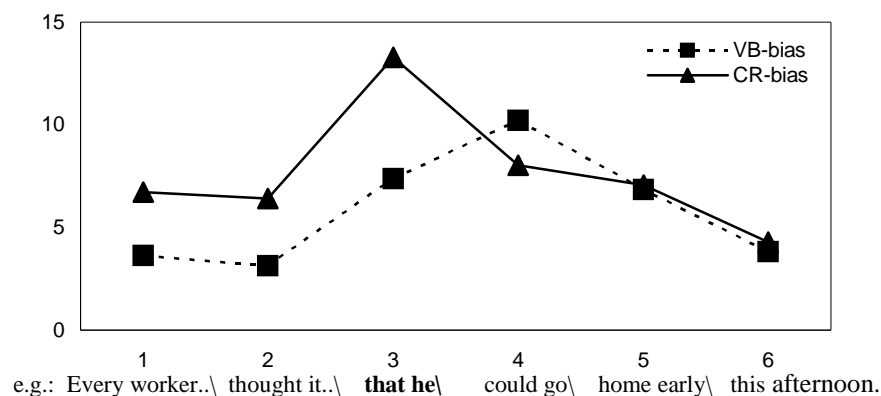
A working day in the factory is always very tough. Especially today the very old man Paul could barely cope. *Every worker who knew that Paul*_i was running out of energy, thought it was very nice that **he**_i could go home early this afternoon. After a hot shower things would probably look better.

Whereas readers did not seem to have obvious processing difficulties during the initial interpretation of the CR-biased pronoun as indicated by the absence of significant first-pass differences in any of the regions, they clearly slowed down during later stages of processing, as significant longer reading times emerged in the CR-biased condition in the second-pass measure (see Figure 1). More specifically, readers re-fixated the critical region (i.e., containing the pronoun) and the preceding region longer in the CR-biased condition than in the VB-biased condition. This finding, particularly the difference in the pre-critical region, suggests that readers were engaged in re-analysis when they read a CR-biased

³ We tested our materials in two internet-based questionnaires. In one of them we simply asked participants to indicate how they interpreted the pronoun by giving them a 2-choice option between the variable binding antecedent and the co-reference antecedent. In the other we presented the stories up to the critical pronoun and instructed the participants to finish the story with the first ending that came to mind. Both experiments revealed an equally strong bias towards the intended interpretation.

context. Hence, this result is consistent with the idea that a variable binding reading is initially preferred over a co-reference reading even if the context strongly supports the latter.

Figure 1. Mean Second-Pass Durations in msec per character for the Variable Binding Biased (VB-bias) and Co-reference Biased condition (CR-bias).



To test whether this preference for a variable binding interpretation occurs because logical syntactic computations are intrinsically less costly than discourse computations, we presented stories containing unambiguous pronouns like (8) and (9). In both stories the pronoun ‘he’ has only one possible referent (i.e., ‘every worker’ or ‘Paul’). By hypothesis, the process that assigns the referent to the pronoun differs between conditions. In (8) the pronoun can only be resolved at the logical syntactic level because the antecedent is quantified, and therefore, by definition not referential. On the other hand, in (9) no c-command relation exists between ‘Paul’ and ‘he’ and the pronoun can only receive its value through co-reference. If co-reference is intrinsically more costly than variable binding we should expect longer reading times in (9) around the critical pronoun.

Furthermore, we included stories like (10) and (11) to check whether the original or revised version of Rule I is preferred. According to the original version of Rule I, a costly comparison between variable binding and co-reference should be made in (10) and (11) because both processes are available in principle. Consequently, longer reading times are predicted in story (10) and (11) than in (8) and (9). In contrast, revised Rule I does not predict a difference in reading times as this version suggests that the rule is only executed if an antecedent in a potential binding position is ruled out by the grammar, which is not the case in any of the stories.

(8) Variable binding only (VB-only)

A working day in the factory is always very tough. Especially today a lot of workers could barely cope. **Every worker_i** who was running out of energy, thought it was very nice that **he_i** could go home early this afternoon. After a hot shower things would probably look better.

(9) Co-reference only (CR-only)


A working day in the factory is always very tough. Especially today the very old man Paul could barely cope. **Paul_i** was running out of energy. It was very nice that **he_i** could go home early this afternoon. After a hot shower things would probably look better.

- (10) Variable binding/co-reference 1 (VB/CR1)
 A working day in the factory is always very tough. Especially today a lot of workers could barely cope. **All workers_i**, who were running out of energy, thought it was very nice that **they_i** could go home early this afternoon. After a hot shower things would probably look better.
- (11) Variable binding/co-reference 2 (VB/CR2)
 A working day in the factory is always very tough. Especially today the very old man Paul could barely cope. **Paul_i**, who was running out of energy, thought it was very nice that **he_i** could go home early this afternoon. After a hot shower things would probably look better.

We observed no significant differences in first-pass, regression path and total reading durations in any of the regions. The analysis of the second-pass durations, on the other hand, revealed an unexpected effect. In the Variable Binding-only condition longer second pass reading times emerged in the region that directly followed the critical pronoun (region 4, see Figure 2). The comparisons between the other conditions revealed no significant differences. As such, this finding provides no support for the hypothesis that discourse operations require more processing resources than logical syntactic operations.

To examine the predictions of the different versions of Rule I, we compared the mean reading times collapsed over the conditions in which only one route was available (the *only* conditions) to the mean reading times collapsed over the conditions in which both routes were available (the *VB/CR* conditions). The analyses revealed no significant difference in any of the regions. These results are consistent with the predictions of revised Rule I, and not with original Rule I.

Figure 2. Mean Second-Pass Durations in msec per character for the Variable Binding Only (VB-only), Co-reference Only (CR-only) and Variable Binding/Co-reference Conditions (VB/CR1 and VB/CR2).



Conclusions

In an eye-tracking experiment we manipulated the interpretation of ambiguous and unambiguous pronouns in an attempt to explore how variable binding, co-reference and Rule I influence the way readers resolve pronouns. The results for unambiguous pronouns

revealed that no processing differences existed for pronouns that could be connected to their antecedents by only one route or, alternatively, by both routes. This result seems inconsistent with the original idea that Rule I always compares the variable binding interpretation to the co-reference interpretation and opts for the former unless the two potential dependencies yield different interpretations (Grodzinsky & Reinhart, 1993). The absence of a costly (and intuitively redundant) comparison between the two derivations in sentences with grammatical variable binding dependencies is, on the other hand, completely consistent with the predictions of revised Rule I, in which it is assumed that the rule is only relevant under specific circumstances (Reinhart, 2000). That is, the parser does not seem to apply Rule I blindly, but instead only uses it for structures similar to *the clown_i hates him_i*. These structures never allow a variable binding interpretation as the pronoun is structurally too close to the antecedent (i.e., a violation of Principle B), yet presented within the right context a co-referential dependency can be established, provided that revised Rule I allows it (see Footnote 2 for an example).

Consistent with Reuland's (2001) processing hierarchy for anaphor resolution we found that ambiguous pronouns are preferably connected to a variable binding antecedent even if discourse information clearly favoured a competing co-reference antecedent. This suggests that the preference for variable binding dependencies is not restricted to VP-ellipsis, but is a general tendency of the human parser during anaphor resolution. Furthermore, as discourse constraints seem to be unable to immediately influence the preference of the parser, we propose that the variable binding interpretation is established independently from the discourse storage, which is consistent with a two-route architecture for pronominal resolution (Reinhart, 1983; Grodzinsky & Reinhart, 1993; Reuland, 2001). Moreover, our results indicate that it may be warranted to consider logical syntax and discourse as distinct modules of the language system (Reuland, 2001).

As noted, if the construction of a dependency in the logical syntactic module has precedence over the construction of a dependency in the discourse module, this can either indicate that the former is intrinsically more economic than the latter, or it can be a time course effect, in the sense that the logical syntax module is accessed earlier. For the latter issue it may then be relevant that, in the case of unambiguous pronouns we found that a pronoun that could only be resolved through variable binding required *more* processing resources than a pronoun that could only be resolved through the co-reference route. This result appears to go against an approach based on intrinsic economy. It could be caused by the specific antecedents we used in our stories.⁴ Burkhardt (2004), for instance, suggests that quantified antecedents like '*every worker*' are more complex than non-quantified antecedents like '*the worker*' or '*Paul*' and, as a result, increase the processing load during pronoun resolution. In her account the most economic dependency between a pronoun and antecedent is established with so-called 'light' quantifiers like '*everyone*'.⁵ If it is indeed more costly for the parser to connect a pronoun to an antecedent like '*every worker*' than to an antecedent like '*Paul*', the preference for the quantified antecedents in our ambiguous materials cannot be exclusively attributed to a difference in processing costs for variable

⁴ As one anonymous reviewer pointed out, the critical pronouns in the CR-only and VB-only conditions appeared in different structural contexts, which may have caused the unexpected effect. However, this would left unexplained why the reading times for the CR-only condition did not differ from the VB/CR conditions as in the latter conditions the structural context was exactly the same as in the VB-only condition.

⁵ Currently we are conducting an eye-tracking experiment in Dutch in which we compare the interpretation of pronouns with light-quantified, quantified and non-quantified antecedents.

binding and co-reference. This might suggest that the language system prefers a binding interpretation, not because it is by definition more economic (that depends on the type of the antecedent), but due to the time course of processing. That is, the parser initially considers c-commanding antecedents by consulting the logical syntactic module and only later has access to the discourse module after which other (sometimes more suitable) antecedents are evaluated. At this point, however, this explanation is rather speculative and needs to be examined more explicitly in future experiments.

References

Arnold, J., Eisenband, J. G., Brown-Schmidt, S., & Trueswell, J.C. (2000). The rapid use of gender information: Evidence of the time course of pronoun resolution from eye tracking. *Cognition*, 76, B13-B26.

Burkhardt, P. (2004). *Representation and Interpretation at the Syntax-Discourse Interface: Establishing Dependency*. Ph.D. Dissertation, Yale University, New Haven, USA.

Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behavior*, 16, 601-609.

Clark, H. H., and Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory and Cognition*, 7, 35-41.

Ehrlich, K. (1980). Comprehension of pronouns. *The Quarterly Journal of Experimental Psychology*, 32, 247-255.

Frazier, L., & Clifton, C. Jr. (2000). On bound variable interpretations: the LF-Only hypothesis. *Journal of Psycholinguistic Research* 29, 125-139.

Greene, S. B., McKoon, G., & Ratcliff, R. (1992). Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 182, 266-282.

Grodzinsky, Y., & Reinhart, T. (1993). The innateness of binding and coreference. *Linguistic Inquiry*, 24, 69-101.

Koornneef, A.W. & Van Berkum, J.J.A. (2006) On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language*, 54(4), 445-465.

McDonald, J. L. & MacWhinney, B. (1995). The time course of anaphor resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34, 543-566.

Reinhart, T. (1983). *Anaphora and Semantic Interpretation*. University of Chicago Press, USA.

Reinhart, T. (2000). Strategies of anaphora resolution. In *Interface Strategies*, ed. Bennis, Everaert, & Reuland, 295-325. Amsterdam: Royal Academy of Arts and Sciences.

Reuland, E (2001). Primitives of binding. *Linguistic Inquiry* 32, 439-492.