

Processing Evidence for Quantifier Raising: the Case of Antecedent Contained Deletion

Jorie Koster-Moeller, Jason Varvoutis, and Martin Hackl
Pomona College

1. Introduction

This paper argues that two combinatorial processes, resolving antecedent contained deletion and resolving the problem of quantifiers in object position, are intimately linked and then discusses the theoretical implications that this hypothesis carries. The evidence that we use to support this claim comes from a series of sentence processing studies that investigate real time effects of processing quantificational and definite DPs.

It is well known that quantificational and definite DPs exhibit essentially the same distributional properties in surface syntax.¹ This is somewhat surprising since their semantics differs notably. Specifically, quantificational DPs do not refer, and so, unlike definite DPs, their semantics cannot be given in the form of singular terms. This notwithstanding, they appear to combine freely with expressions denoting predicates of individuals. This fact indicates that the language faculty possesses a considerable degree of combinatorial flexibility. What the nature of this flexibility is and whether it should be located in the computational system itself (including lexically encoded combinatorial properties) or elsewhere is the topic of an ongoing debate.²

The literature on this topic can be divided into two main camps. On the one hand, there are theories in which the computational engine is insensitive to the referential vs. non-referential distinction (e.g. Kempson et al. 2001). Theories of this sort straightforwardly predict the syntactic distribution of quantificational DPs to be the same as those of referring DPs, but might require additional (post-syntactic and post-compositional-semantic) machinery to derive the entailment patterns that typically distinguish quantifiers from referential expressions. On the other hand, theories within the neo-Montagovian tradition maintain that the computational engine is sensitive to the distinction and employs mechanisms that affect either the compositional semantics of quantifiers, the syntax of quantifiers, or both in order to guarantee interpretability in all their syntactic occurrences. These theories can thus be further distinguished depending on whether these mechanisms are semantic (e.g. type-shifting) or syntactic (e.g. Quantifier Raising).

In previous work, Varvoutis & Hackl (2006) presented experimental evidence showing that the semantic status of a DP, whether it is quantificational or not, is a factor in local ambiguity resolution during first pass parsing. More specifically, they show that quantifiers are easier to interpret in real time when they are in subject position than when they are in object position.³ These results support the view that the computational system does distinguish definite descriptions from quantificational DPs, but do not allow us to choose between the

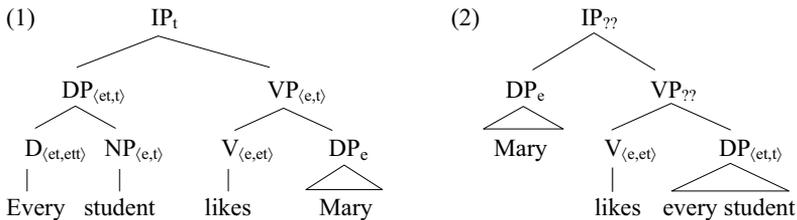
various approaches to the syntax-semantic mapping problem that arises within the neo-Montagovian camp.

The goal of the current paper is to provide experimental evidence that does allow us to choose between different ways of encoding the increased complexity that quantifiers in object position pose.

2. Quantifiers in object position

It is well known that the assumption that the computational system distinguishes between referring and non-referring expressions implies that quantifiers in object position present a case of syntax-semantics mismatch.⁴

Assuming familiar type theory to characterize the combinatorial properties of natural language expressions, the mismatch arises as follows: Predicates such as *like*, given that their conceptual core is that of a relation between individuals, are functions from individuals to sets of individuals (type $\langle e, et \rangle$).⁵ This allows them to combine directly with referring expressions (type e) such as proper names and definite descriptions in both subject and object position. Quantifiers, however, given that they do not denote individuals, are not of type e . Instead, they are standardly viewed as second order predicates (type $\langle et, t \rangle$). This introduces a subject-object asymmetry since quantifiers can combine directly with their sister when they occur in subject position, (1), but not when they occur in object position, (2). The reason is that expressions of type $\langle e, et \rangle$ cannot take a quantifier, type $\langle et, t \rangle$, as argument nor the other way around.



Various solutions to the problem of quantifiers in object position (QOb) have been proposed in the literature. Below we sketch two approaches to illustrate a difference between semantic and syntactic solutions that is relevant for our experimental research.

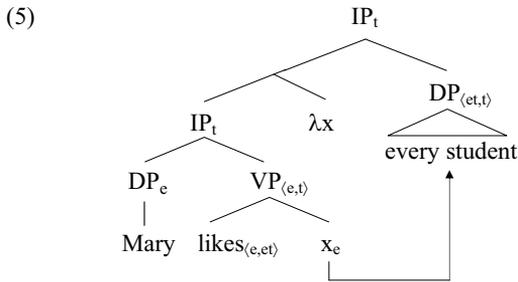
One way to address QOb is to assume that transitive verbs have a secondary meaning that allows them to take expressions of type $\langle et, t \rangle$ as arguments, (3).⁶ The secondary meaning can be derived from the primary meaning, (3), by applying a type-shifting operation, (4), which can be thought of as either a lexical rule or a rule of semantic composition.⁷

- (3)
- a. $\llbracket \text{likes}_1 \rrbracket = \lambda x: x \in D_e. \lambda y: y \in D_e. y \text{ likes } x.$
 - b. $\llbracket \text{likes}_2 \rrbracket = \lambda F: F \in D_{ett}. \lambda x: x \in D_e. F(\lambda y. \llbracket \text{likes}_1 \rrbracket(y)(x)=1)=1.$

$$(4) \quad \text{TSH} = \lambda f: f \in D_{\text{et}}. \lambda F: F \in D_{\text{et}}. \lambda x: x \in D_e. F(\lambda y. f(y)(x)=1)=1.$$

Importantly, type-shifting solutions to QOb such as the one sketched above do not involve any syntactic reorganization of the clause. The quantifier *every student* can be composed in situ by employing a more complex variant of the basic transitive verb. This is seen as a conceptual advantage in Jacobson (1996, to appear) as well as in Barker's (2001) continuations approach to QOb.

Syntactic alternatives to purely semantic approaches to QOb assume that the type-mismatch can be resolved by dislocating the quantifier to a position where it can be interpreted as a variable binding operator that binds a variable inserted in the quantifier's original position.⁸ The best known version of this type of approach assumes that the quantifier is moved covertly by a rule called Quantifier Raising (QR) to a clause initial position (cf. May 1977) as sketched in (5).



We assume in (5), following Fox (2002, 2003), that QR is an instance of covert rightward movement, which adjoins the object quantifier *every student* to a clause initial position leaving behind a co-indexed trace of type e .⁹ The trace is interpreted as a variable bound by a λ -operator, which is introduced as the moved QP is adjoined to IP.

Choosing between different approaches to QOb based on off-line data alone is a complex task because it requires comparisons of entire frameworks.¹⁰ Our contribution to this debate comes instead from real time sentence processing, which offers a method for investigating QOb that is not available for traditional off-line investigations.

Before we describe our experimental work, we review briefly the phenomenon of antecedent contained deletion, which is central to the debate between syntactic and semantic approaches to QOb, as well as the basis of our experimental paradigm.

3. Antecedent contained deletion

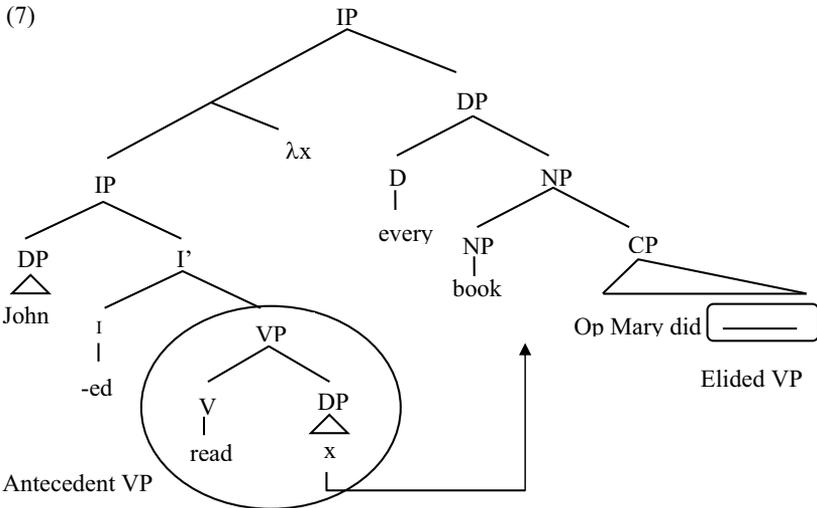
The term antecedent contained deletion (ACD) refers to elided material, marked by in our examples, that is properly contained within the expression that serves as its antecedent, (6).

(6) John read every book Mary did .

Sentences like (6) are standardly analyzed as cases of VP ellipsis inside the relative clause. Importantly, the antecedent of the elided VP is the matrix VP, which contains the object DP that hosts the ellipsis site.

The existence of ACD is puzzling because ellipsis is generally subject to a licensing condition that demands the elided material to be identical (parallel) to a pronounced antecedent constituent (cf. Sag 1976 etc.).¹¹ Since the ellipsis site in ACD configurations is properly contained within the expression that serves as its antecedent, the identity requirement seems impossible to satisfy. In (6), the purported antecedent of the elide VP inside the relative clause is the matrix VP *read every book Mary did* , which is clearly not identical to .

As before, it is possible to distinguish two types of approaches to the paradox presented by ACD. One approach assumes that ACD involves covert movement of the DP hosting the ellipsis site out of the antecedent constituent, (7).¹²



As is evident from (7), movement to resolve antecedent containment can be seen as an instance of QR since it involves covert phrasal movement of the object DP to a clause initial position. Indeed, ACD resolution shares many of the

properties of QR and has consequently been argued to provide independent evidence for its existence (May 1985, Heim & Kratzer 1998, etc.).

An alternative approach to ACD, defended in Jacobson (to appear), is to assume that ACD involves V-ellipsis where the elided material is simply a transitive verb rather than VP ellipsis. According to this analysis, there is no antecedent containment and consequently no need to move the object DP to undo containment.¹³

A detailed comparison of these two types of approaches is beyond the scope of this paper. What is important for our purpose here is the fact that the two approaches differ in the implications that ACD has for the host DP.

The first approach, which assumes antecedent containment and covert phrasal movement as the mechanism that undo it, predicts that any DP, whether quantificational or not, will undergo QR if it hosts a relative clause containing an ACD site. The V-ellipsis approach to ACD, on the other hand, does not make that prediction.

There is off-line support for a QR based view of ACD. The most prominent type of support comes from the fact that it predicts the Sag-Williams generalization, according to which the semantic scope of the object DP hosting the ACD site is at least as high as the antecedent constituent. Consider the sentences in (8) after Sag (1976).

- (8) a. Betsy's father wants her to read everything her boss does.
 b. Betsy's father wants her to read everything her boss wants her to read.

(8) is ambiguous. It can be paraphrased either by "every x is such that if x is something that Betsy's boss wants her to read then x is something that Betsy's father wants her to read it" (*de re*) or by "What Betsy's father wants is that she read everything that her boss wants her to read" (*de dicto*). Interestingly, (8) does not have the second reading.¹⁴ Sag suggests that the absence of that reading should be related to the fact the object DP needs to out-scope *want* if the antecedent of the elided VP is to be understood as "wants her to read." Since the object DP is higher than *want* under that ellipsis resolution it cannot bind any world variables inside the object DP and this explains that absence of the *de dicto* reading.

Sag's account can be adopted by a QR approach to ACD, which simply adds a syntactic mechanism that guarantees the relevant scope relation. Specifically, a QR approach predicts that the object DP moves higher than the antecedent VP in order to undo antecedent containment. If the antecedent is the matrix VP the object DP has to out-scope the modal *want*, which is sufficient to prevent a *de dicto* construal of the object DP.

Note that the reading is absent whether the object DP is quantificational or not. Specifically, the variant in (9) displays the same effect. It cannot be paraphrased by "What Betsy's father wants is that she read the same book that here boos wants her to read."

- (9) Betsy's father wants her to read the (same) book that her boss does.

This is predicted under the QR approach to ACD since QR of the object DP over the matrix predicate *want* is necessary as soon as the antecedent for ellipsis resolution includes *want*. Thus, the object DP is moved over the matrix verb no matter what its semantic properties are. Specifically, the object DP does not need to be quantificational to undergo QR in this case.

It is less obvious to us that V-ellipsis approaches to ACD, such as the one defended in Jacobson (to appear), predict the full correlation discussed above. According to Jacobson, in (8) the antecedent of the elided constituent is the complex verb *wants her to read*. Because the licensing condition on ellipsis requires semantic identity (parallelism) the complex verb needs to be assembled compositionally (via function composition) before it can serve as an antecedent. Once the assembly is completed, the result is a transitive verb. Whether it takes scope over the object depends on the semantics of the object DP. For quantificational objects, Jacobson assumes a type-shifted version of the quantifier that takes the transitive verb as argument. This is sufficient to disallow a *de dicto* reading since the modal operator that is part of the complex verb *wants her to read* is effectively out-scoped by the object DP. If, however, the object DP is a definite DP, as in (9), the prediction changes.

Since transitive verbs can take definite DPs as arguments, the verb has scope over the DP. The availability of a *de dicto* reading is therefore contingent upon how the complex verb is assembled. Specifically, if the modal operator inside the complex verb can bind a world variable associated with the DP object, a *de dicto* reading should be possible. This is predicted under Jacobson's proposal because there is no principled reason why the complex verb couldn't be opaque (taking an argument of type $\langle s,e \rangle$) with respect to the object position. It seems to us, however, that a *de dicto* reading for (9) is as unavailable as it is for (8).

We leave a more detailed discussion of this issue for another occasion and instead focus on a different way of examining whether or not semantic properties of the object DP are systematically linked with ACD resolution. We will show that real time sentence processing can be used to construct a new type of argument in favor of a QR based account of QOb and ACD.

4. Processing antecedent contained deletion

Recall that QR based accounts of QOb and ACD assume that there are two types of triggers for QR: the presence of a quantifier in object position and an ellipsis site that is properly contained within its antecedent. Importantly, these two triggers are independent of each other. Thus, if the host DP of a potential ACD site is quantificational, QR is triggered whether the relative clause has a pronounced or elided VP. If, however, the object DP is referential, QR is triggered only if the relative clause contains an ACD site. In other words, QR based approaches to QOb and ACD predict an interaction between DP type and presence/absence of ACD such that QR is triggered in all cases except when the object DP is referential and does not host an ACD site.

Since QR is a covert movement operation that has no phonetic reflex, constructing off-line data that would constitute compelling evidence for such an interaction is difficult. Real time sentence processing offers a distinct advantage here because the two triggers occupy different positions in the string (the object DP precedes the ellipsis site)¹⁵ and sentence processing is inherently sensitive to linear order.

To see how sentence processing might reveal an interaction between determiner type and presence/absence of an ACD site (“gap-size”), consider what a QR approach would predict about processing difficulty for reading material after the gap, which is a simple trace in (10) and an ACD site in (10), in sentences such as those in (10) (QR triggers are marked in bold face).

- (10) a. John talked to the student that Mary talked to before class.
 b. John talked to the student that Mary **did** before class.
 c. John talked to **every** student that Mary talked to before class.
 d. John talked to **every** student that Mary **did** before class.

In (10) the object DP is a definite DP, which does not trigger QR. In (10) the gap inside the relative clause is a simple trace rather than an elided VP, so there is no trigger for QR inside the relative clause, and QR should not occur at all. (10), on the other hand, triggers QR at the ellipsis site. Assuming that QR and ellipsis resolution incur a processing cost, we would expect increased processing difficulty for (10) at or after *did*. (10) are constructed identical to (10) except that the definite determiner is replaced by *every*. Since this makes the object DP quantificational, QR is triggered at the location of the quantificational determiner in both cases. Importantly, since QR has already been triggered by *every*, no additional QR-related processing cost should arise at the ellipsis site in (10) compared to (10). In other words, the QR approach to QOb and ACD predicts an increase in complexity between trace and ellipsis when the host DP is a definite description but no comparable increase when the host DP is quantificational. Since QR has already been triggered earlier, the parser is fully prepared to accommodate elision of the VP.

The current experiment follows the same logic with one addition. Our experimental paradigm, exemplified in (11), includes two types of ACD sites, a small ellipsis site whose antecedent is the embedded VP and a large ellipsis site whose antecedent is the matrix VP.

- (11) The secretary was trained to manage...
- a. the/every program that the intelligent young professional designed ...
 b. the/every program that the intelligent young professional did ...
 c. the/every program that the intelligent young professional was ...

...during her four years at college.

As discussed above, the QR approach to QOb and ACD predicts an interaction between determiner type and gap size for (11) at or after the gap.

Predictions regarding an increase in processing difficulty for the quantificational variant of (11) (*every program*) depend on whether QR is assumed to be, by default, local. If QR is assumed to be local, *was* should trigger a second application of QR and therefore an increase in processing cost as compared to (11). The definite variant of (11) (*the program*) provides a baseline that includes long QR and large ellipsis resolution.¹⁶

4.1 Methods and Materials

To investigate whether real time processing of ACD partially depends on the semantics of the host DP as discussed above, we use the self-paced, word-by-word moving window reading methodology (Just, Carpenter, & Woolley 1982).

Our target items were constructed following the sample paradigm in (11). The matrix verb was always in the past progressive to allow for large ellipsis resolution triggered by *was* in the relative clause. Embedded verbs were chosen so that the progressive reading (*was managing*) was strongly dispreferred.¹⁷

Adverbs and adjectives were inserted between the object DP and the main point of interest (the verb or auxiliary in the relative clause) to prevent spillover effects from the different determiners to interfere with processing difficulties that might arise at the point of interest.

We constructed 60 target sentences, which were combined with 120 fillers of various types. These included sentences that were similar to the target items in structure (employing relative clauses, elided material or covert movement triggers), in length, or because they contained quantifiers. The items were counterbalanced across six lists using a Latin-square design.

Forty-eight undergraduates from the Claremont colleges were tested on a Dell computer running the Linger software developed by Doug Rohde. All were native speakers of English and received course credit or \$10.00 cash for their participation. Each trial begins with a series of dashes marking the length of the sentence. Participants press the spacebar to reveal the next word of the sentence. Each press of the space bar reveals a new word while the previous word is again replaced by dashes. The amount of time a participant spends reading each word is recorded (RT). After the final word of each sentence, a yes/no comprehension question appears, asking about information contained in the sentence. Participants respond by pressing keys marked on the keyboard. If an incorrect answer is given, the word "incorrect" appears on the screen to encourage participants to read more carefully. Participants are instructed to read sentences at a natural rate and to be sure that they understand what they read. They are also instructed to answer the questions as quickly and accurately as possible.

Items were pseudo-randomized separately for each participant, with at least one filler sentence preceding each target. Before the experiment began, a set of practice items was presented. Participants took approximately 45 minutes to complete the experiment. Participants were instructed to take short breaks during the experiment to prevent effects of fatigue or habituation to the task.

4.2 Results

Questions for experimental items were answered correctly on 85% of trials. Accuracy rates were 85% for both the definite and quantificational determiner conditions. Accuracy rates were 83.6% for basic verb sentences, 85.5% for small ellipsis sentences (*did*) and 85.8% for sentences containing a large ellipsis site (*was*). For the verb conditions, accuracy rates were 84% for definite determiners and 83.1% for quantificational determiners. For the small ellipsis conditions, accuracy rates were 84.8% for definite determiners and 86.1% for quantificational determiners. For the large ellipsis conditions, accuracy rates were 85.7% for definite determiners and 85.9% for quantificational determiners. A two-factor ANOVA crossing determiner (definite versus quantificational) with ellipsis-type (none, small, large) on the question-answering data revealed no significant differences.

To adjust for differences in word length and differences in participants' natural reading rates, a regression equation predicting reading time from word length was derived for each participant, using all items. At each word position, the reading time predicted by the participant's regression equation was subtracted from the actual measured reading time to obtain a residual reading time. Residual reading times beyond 2 standard deviations from the mean for a given condition and position were excluded from analyses. Figure 1 shows residual reading times by condition.

A repeated measures ANOVA on the two words following the determiner (*program that*) reveals a main effect of determiner such that quantificational DPs are read more slowly than definite DPs [$F(1,47) = 4.591, p = 0.037$].

At the region of interest, two words after the ellipsis site (marked by *her* in Figure 1), we see a prominent separation of reading times across conditions.¹⁸ A repeated measures ANOVA (Determiner by Gap size) reveals a main effect of Gap size [$F(2,46) = 12.969, p < .001$] and a significant interaction [$F(2,46) = 4.363, p = 0.18$].

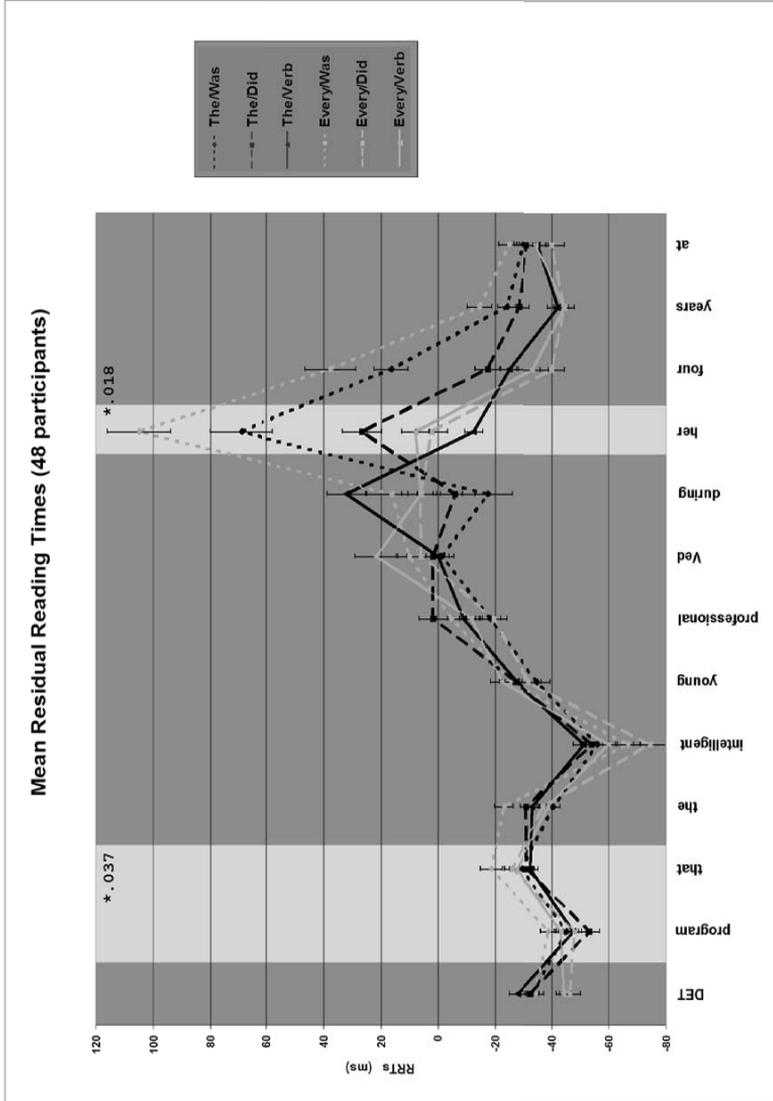


Figure 1. Mean residual reading times across 6 conditions, (48 subjects).

The interaction is driven by the fact that the small ellipsis site in the case of a quantificational host DP (*every-did*) does not display the same linear increase in reading time across the three gap sizes we see with definite host DPs. Specifically, in the definite determiner case there is a linear effect of gap size ($V < did < was$) while in the *every* case there is no difference between *every-V* and *every-did* ($V, did < was$). This can be seen clearly in Figure 2, which presents the three different levels of gap size (*Verb*, *did*, and *was*) on the horizontal axis.

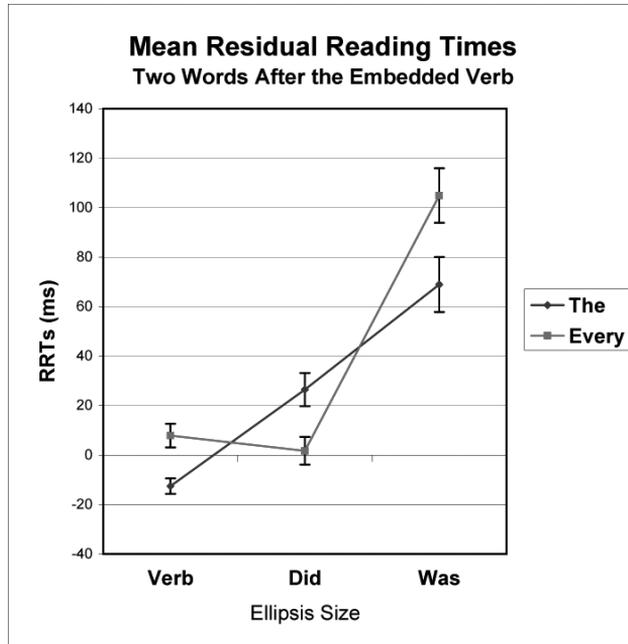


Figure 2. Mean residual reading times 2 words after V/Aux, (48 subjects).

In sentences with quantificational determiners, large ellipsis versions were read more slowly than versions with either small or no ellipsis [$F(1,47) = 21.329$, $p < .001$]. There was, however, no significant difference between small and no ellipsis. This contrasts markedly with reading times for the definite determiner condition.

In fact, a comparison of the definite determiner condition and the quantificational condition across the small and no ellipsis conditions reveals a significant interaction of determiner type by gap type [$F(1,47) = 4.611$, $p = .037$]. A comparison of small and large ellipsis across the two determiner types reveals a main effect of gap size, [$F(1,47) = 17.074$, $p < .001$], and a significant interaction, [$F(1,47) = 6.125$, $p = .017$], which is again generated by the fact that reading times for *every-did* are comparatively lower than reading times for *the-did*.

4.2 Discussion

The main effect of determiner observed two words after the determiner (*every NP* > *the NP*) replicates a parallel finding reported in Varvoutis & Hackl (2006) showing that participants react to the difference between *the* and *every*. The fact that *every NP* is harder to process supports the idea that quantificational DPs are more complex for the parser than definite DPs. It doesn't, however, show whether this difference is due to a difference in combinatorial complexity (quantifiers are of type <et,t> while definite DPs are of type e) or a difference in conceptual semantic complexity (the denotation of *every NP* is arguably more abstract than the denotation of *the NP*).¹⁹ To argue convincingly that the parser integrates quantifiers differently than it does definite DPs, the evidence needs to come in the form of an interaction that takes a structural property of the sentence into account.

This type of evidence can be seen in the interactions we observe two words after the ellipsis site, which provide compelling support for the hypothesis that the semantic property of the host DP affects ACD resolution.

The interaction between determiner type and small/no ellipsis is due to a marked increase in reading times between *the-V* and *the-did* that is not found for *every-V* compared to *every-did*.²⁰ The interaction between determiner type and small/large ellipsis, on the other hand, is due to a larger increase in reading time for *every-was* relative to *every-did* compared to the increase in reading time for *the-was* relative to *the-did*. This indicates that the semantic property of the host DP affects real time processing of ACD and, to the extent that first pass parsing is sensitive to structural differences, that quantifiers and definite descriptions are integrated differently. Theories that do not provide the means to distinguish between the combinatorial properties of definite and quantificational DPs are therefore inconsistent with our data.

To see how our data distinguish between various accounts of QOb and ACD, consider first how a QR based approach would explain these interactions. Under the QR approach to ACD, the increase in reading times for *the-did* is due to the cost of QRing the host DP and the cost of accessing the antecedent for *did*.²¹ In the case of *every-did*, QR of the host DP has already been triggered much earlier in the sentence. Therefore, there is no QR-related cost for ellipsis resolution after *did* and consequently less of an increase (in fact no significant increase) in reading times for *every-did* compared to *every-V*.²² Thus, the type of interaction we observe between determiner type and small/no ellipsis is precisely what a QR based approach to QOb and ACD would expect.

The interaction we observe between determiner type and small/large ellipsis might be understood in the same way as the interaction between determiner type and small/no ellipsis, namely the unexpectedly fast reading times for *every-did*. I.e. instead of observing a main effect of determiner across all three gap sizes (*every-V/did/was* > *the-V/did/was*) presumably due to quantificational DPs being more difficult to process than definite DPs, we observe an interaction generated by unexpectedly fast reading times for *every-did* given the baseline reading times provided by *the-did*. If this interpretation is correct, a QR approach to QOb and ACD can account for the results in the same way it accounts for the

interaction between determiner type and small/no ellipsis. Alternatively, one might suspect that the second interaction is driven by a disproportionate increase in reading times for *every-was* relative to *every-did*, given the baseline increase from *the-did* to *the-was*.

This interaction, however it is analyzed, could be due to the fact that the initial application of QR triggered by *every NP* is, by default, local. Since local QR, which adjoins the host DP to the embedded IP, is insufficient to license large ellipsis resolution, a second application of QR, which moves the host DP into the matrix, is necessary.

Consider next what semantic approaches to QOb and ACD such as Jacobson (to appear) would predict for our cases. The increase in reading times we observe for *the-did* compared to *the-V* after the ellipsis site can be attributed to ellipsis resolution. Of course, since exactly the same process occurs in the case of *every-did*, Jacobson would predict the same increase for *every-did* compared to *every-V*. Since Jacobson treats ACD as a case of V-ellipsis, the semantics of the object of V cannot have a direct effect on ACD resolution. Hence, this account predicts a main effect of ellipsis when comparing reading times as a function of the two factors *the/every* and *V/did*. Reading times for *did* should be longer than those for *V* and it shouldn't be effected by the type of the determiner (*the/every-did* > *the/every-V*). Of course, this is contrary to fact. The interaction we observe, that a quantificational host DP, compared to the definite host DP, makes ACD resolution easier, cannot be explained in any theory that, like Jacobson's, treats ACD essentially as V-ellipsis, irrespective of what assumptions are made about QOb.^{23,24}

It is worth pointing out that our data are inconsistent with whole classes of ACD and QOb accounts. Specifically, our findings are problematic for all ACD accounts that are not sensitive to the semantic status of the object DP. Hornstein (1995), for instance, assumes that all objects are moved to a functional projection for case reasons (A-movement). This is, as Hornstein points out, high enough to license ACD. Since case driven movement targets quantificational DPs and definite DPs alike, Hornstein would not predict any difference in ACD resolution depending on the semantic status of the host DP.²⁵

Essentially the same argument can be made against theories such as Kempson et al. (2001), which do not distinguish between the combinatorial properties of quantificational and definite DPs in general, hence do not recognize a difference in compositional complexity between quantifiers and definite DPs in object position. Our findings are inconsistent with such theories because treating all DPs uniformly (as expressions of type *e* in Kempson et al 2001) when appearing in object position removes the means for the parser to predict a difference in structure that might become relevant further down stream. Thus, uniform treatments of DPs do not provide the means to link QOb and ACD, which means that they cannot explain how a quantifier in object position could make ACD resolution easier.²⁶

5. Conclusion

In this paper we have presented real time sentence processing evidence that links two compositional processes: the problem of quantifiers in object position (QOb) and antecedent contained deletion (ACD). The argument that these two processes are linked came in the form of an interaction in reading times of sentences that varied across two factors. The first factor concerned the semantic status of the host DP (referential, *the NP* or quantificational, *every NP*). The second factor concerned the gap size inside the relative clause, which varied between a simple trace, a small ellipsis (*did*), whose antecedent was the embedded VP, or a large ellipsis (*was*), whose antecedent was the matrix VP.

The key finding of our study is an interaction in reading times after the gap. Specifically, we observed a linear increase in reading times for definite host DPs across the three gap sizes but not for quantificational host DPs. Rather than a linear increase, we found that there was no difference between sentences that had no ellipsis and sentences that had a small ellipsis site, but a large increase when the ACD site required a matrix antecedent.

The fact that there was no difference between *every-V* and *every-did* strongly suggests that a quantifier in object position facilitates small ACD resolution later in the sentence. This finding is naturally explained by theories that link these two processes and cannot be explained without making ad hoc assumptions by theories that don't. A direct way to link these two processes is to assume, as QR based accounts of QOb and ACD do, that they are executed by the same mechanism. This is probably not the only way to link QOb and ACD. However, to the extent that QR is seen as the only mechanism that does, our paper provides evidence for the existence of QR.

Endnotes

Portions of this paper have been presented at the Research Seminar in Linguistics and Cognitive Science at Pomona College, CUNY 2007, and WCCFL 26. We would like to thank the audiences at these venues for helpful comments and suggestions as well as Chris Barker, Lynn Frazier, and especially Elena Guerzoni.

¹ This is true for languages like English but not for all languages. Cf. Szabolcsi (1997), etc.

² We use the label "computational system" in this paper to refer to the syntactic and compositional semantic component of the grammar.

³ Varvoutis & Hackl (2006, to appear) use sentences such as those in (i) and (ii), which are locally ambiguous between construing the DP *the/every witness* as object of *believe* or as subject of an embedded clause (NP/S ambiguity), to investigate whether first pass parsing is sensitive to the referential/non-referential distinction.

-
- (i) The judge believed the witness ...
 - a. ...was at the scene of the crime.
 - b. ...who was at the scene of the crime.
 - (ii) The judge believed every witness ...
 - a. ...was at the scene of the crime.
 - b. ...who was at the scene of the crime.

A comparison of readings after the point of disambiguation reveals an interaction of determiner type by attachment type such that attaching the quantificational DP *every witness* as object of *believe*, (iib), causes a significant slowdown in reading times compared to the other three conditions. This can only be explained by a theory that able to identify quantifiers in object position as a case that involves an increase in complexity.

⁴ We follow the presentation of the problem of quantifiers in object position given in Heim & Kratzer (1998).

⁵ We will ignore the distinction between characteristic functions and the sets they characterize.

⁶ Alternatively, one could shift the meaning of the quantificational determiner so that it projects a quantifier that can take a transitive verb as argument.

⁷ Montague (1973) among others assumes that transitive verbs lexically have a meaning of type $\langle \text{ett}, \text{et} \rangle$ while Partee & Rooth (1983) assume that type-shifting is a process that the compositional machinery can call on whenever necessary.

⁸ First order predicate logic employs a syntactic solution of this kind, as all quantifiers, whether they are in subject or object position, are represented by variable binding operators that are prefixed to open sentences.

⁹ Rightward movement seems more natural for sentence processing purposes. Note however, that this is assumption is not essential for our paper.

¹⁰ See, among others, Fox (2003) and Jacobson (to appear) for discussion.

¹¹ Sag's (1976) original formulation requires the logical form of the elided VP to be an alphabetic variant of the logical form of the antecedent VP. The precise formulation of the identity/parallelism condition on ellipsis is an issue of continued debate. See Fox (1998) Rooth (1992) among many others for discussion. For our purpose it is sufficient to note that the condition requires semantic identity/parallelism.

¹² Cf. Sag (1976), May (1985), Larson & May (1990), etc.

¹³ See Cormack (1984), among others, for a similar proposal. We will use Jacobson (to appear) in our discussion as representative for this class of analyses of ACD. It will be obvious that the conclusions we draw from our experimental findings will hold for the entire class. See Hornstein (1995) for off-line arguments against a (pseudo-)gapping analysis of ACD.

¹⁴ See Williams (1977) for similar observations based on sentences like *Mary's father told her to work harder than her boss did*.

¹⁵ Strictly speaking, this is incorrect since the relative clause containing the ellipsis site is part of the object DP. What we mean by "object DP" here is the string comprised of the determiner and the NP that is modified by the relative clause.

¹⁶ There are various ways in which one can model a difference in processing cost due to gap size (trace, small ellipsis and large ellipsis), e.g. distance based metrics following Gibson (1998) or content based metrics. See Frazier & Clifton (2005) among others for relevant work.

¹⁷ The verbs we chose for the no ellipsis condition were different from the matrix verb and the embedded verb. This was done because we did not want to introduce an experimental bias for small or large ellipsis resolution and because we found in pilot studies that using the same verb as the antecedent caused a slow down compared to the ellipsis variant.

¹⁸ Our target sentences used different adjunct continuations. The second word after V/Aux, specifically, varied in category.

¹⁹ Varvoutis & Hackl's (2006, in preparation) argument that quantifiers are integrated in the sentence differently than definite DPs is based on an interaction between determiner type and attachment type rather than the main effect replicated here.

²⁰ Post-hoc comparisons reveal no significant differences between *every-V* and *every-did* [$F(1,47) = 2.12; p = .647$]

²¹ Our data does not allow us to quantify how much of the cost is due to ellipsis resolution and how much of it is due to QR. Note, however, that the argument does not depend on knowing that. Even if there were no processing cost to QR one would have to explain why ellipsis resolution gets easier if the host DP is quantificational.

²² The fact that there is no increase for *every-did* compared to *every-V* might suggest, in line with Frazier & Clifton (2005), that ellipsis resolution itself does not incur any cost, at least in this condition. Whether this true or not does not affect our argument as already pointed out in fn. 21.

²³ C.f. Cormack (1984), Lappin (1992) for V-ellipsis approaches to ACD, and, the Bachrach & Katzir(2007) for a VP-sharing approach that seems to have the same weakness as V-ellipsis approaches in explaining our findings.

²⁴ In Jacobson's system, in particular, the only feasible way of explaining how a quantificational host DP might affect V-ellipsis resolution would be to assume that the antecedent verb is type shifted to accommodate a quantifier in object position. However, the type-shifted verb will arguably be more fragile, hence harder to maintain in memory than the basic verb, which should mean that ACD should be harder with a quantificational host DP.

²⁵ See Kennedy (1997) for off-line arguments against Hornstein(1995).

²⁶ A similar concern applies to Barker's (2002) continuation approach to QOb, since it does not offer an account ACD. It is less clear whether Montague's (1973) account of QOb would be inconsistent with our data since he assumed a type shifting mechanism as well as a quantifying-in mechanism to deal with quantifiers. However, since Montague assumed that transitive verbs always had a higher type and that all DPs were of type $\langle et, t \rangle$, generalizing to the worst case, we would argue that Montague couldn't account for our data without ad hoc assumption linking QOb with ACD. (Montague did not discuss ACD, of course.)

References

- Baltin, Mark: 1987, 'Do Antecedent Contained Deletions Exist?', *Linguistic Inquiry* **18**.
- Barker, Chris: 2002, 'Continuations and the nature of quantification', *Natural Language Semantics* **10**, 211-242.
- Cormack, Annabel: 1984, 'VP Anaphora: Variables and Scope', in Fred Landman and Frank Veltman (eds.) *Varieties of Formal Semantics*, Dordrecht, Reidel.
- Dalrymple, Mary, Stuart M. Shieber, and Fernando C. N. Pereira: 1991, 'Ellipsis and Higher-Order Unification', *Linguistics and Philosophy* **14**.
- Fox, Danny: 2002, 'Antecedent Contained Deletion and the Copy Theory of Movement', *Linguistic Inquiry* **33**, **1**, 63-96.
- Fox, Danny: 2003, 'On Logical Form', in Randall Hendrick (ed.) *Minimalist Syntax*, Blackwell Publishing Inc.
- Gibson, Edward: 1998, 'Linguistic complexity: locality of syntactic dependencies', *Cognition* **68**:1-76.
- Frazier, Lynn, and Charles Clifton: 2005, 'The Syntax-Discourse Divide: Processing Ellipsis', *Syntax* **8.2**, 121-174.
- Heim, Irene, and Angelika Kratzer: 1998, *Semantics in Generative Grammar*, Oxford, Blackwell.
- Jacobson, Pauline: 1992, 'Antecedent Contained Deletion in a Variable Free Semantics', in *Proceedings of the Second Conference on Semantics and Linguistic Theory*, Ohio State University, 193-213.
- Jacobson, Pauline: 1996, 'Semantics in Categorical Grammar', in Shalom Lappin (ed.) *Handbook of Semantics*, Blackwell Publishing Inc., 89-116.
- Jacobson, Pauline: (to appear), 'Direct Compositionality and Variable-Free Semantics: The Case of Antecedent Contained Deletion', to appear in Kyle Johnson (ed.) *Ellipsis*, Oxford University Press.
- Just, Marcel A., Patricia A. Carpenter, and Jacqueline D. Woolley: 1982, 'Paradigms and processes and in reading comprehension', *Journal of Experimental Psychology: General*, **3**, 228-238.
- Kennedy, Christopher: 1997, 'Antecedent contained deletion and the syntax of quantification', *Linguistic Inquiry* **28.4**.
- Larson, Richard and Robert May: 1990, 'Antecedent containment or vacuous movement: Reply to Baltin', *Linguistic Inquiry* **21**.
- May, Robert: 1977, *The Grammar of Quantification*. Doctoral Dissertation, MIT.
- May, Robert: 1985, *Logical Form: Its Structure and Derivation*, MIT Press, Cambridge, MA.
- Montague, Richard: 1973, 'The Proper Treatment of Quantifiers in Ordinary English', in Richmond H. Thomason. (ed.) *Formal Philosophy*, Yale University Press.
- Partee, Barbara, and Mats Rooth: 1983, 'Generalized Conjunction and Type Ambiguity', in Rainer Bauerle et al. (eds.). *Meaning, Use, and the Interpretation of Language*, de Gruyter, Berlin, 361-393.

-
- Rooth, Mats: 1992, 'Ellipsis Redundancy and Reduction Redundancy', in Steve Berman and Arild Hestvik (eds.) *Proceedings of the Stuttgart Ellipsis Workshop, Report Series, Sprachtheoretische Grundlagen für die Computerlinguistik, no 29*, SFB, 340 IBM Heidelberg: 1-26.
- Sag, Ivan: 1976, *Deletion and Logical Form*. Doctoral Dissertation, MIT.
- Szabolcsi, Anna: 1997. *Ways of Scope Taking*. Kluwer, Dordrecht.
- Varvoutis, Jason and Martin Hackl: 2006, 'Parsing Quantifiers in Object Position', talk delivered at the 19th CUNY Conference on Human Sentence Processing, CUNY, NY. <http://www.linguistics.pomona.edu/mhackl/QobCUNY06.pdf>
- Varvoutis, Jason and Martin Hackl: 2007, 'Parsing Quantifiers in Object Position', ms. Pomona College.
- Wilder, Chris: 2003, 'Antecedent containment and ellipsis', in Schwabe, Kerstin and Winkler, Susanne (eds.) *The interfaces: Deriving and interpreting omitted structures*, John Benjamins, Amsterdam.
- Williams, Edwin: 1974, *Rule Ordering in Syntax*. Doctoral Dissertation, MIT.
- Williams, Edwin: 1977, 'Discourse and logical form', *Linguistic Inquiry* 8, 101-139.