

Discourse Binding: Does it Begin with Nominal Ellipsis?

Frank Wijnen¹, Tom Roeper², and Hiske van der Meulen¹
¹Utrecht University and ²University of Massachusetts, Amherst

1. Introduction

Sentences are normally not used in isolation. They are preceded and followed by other sentences, delivered either by the same speaker or writer, or by an interlocutor. Language users live by the principle that sentences in sequences belong together; that they are meaningfully connected. Language users weld sentences together into a coherent representation of discourse. Coherence is a cognitive state; it is not ‘in’ the language, although sentences may contain devices that help the perceiver in creating it. Sanders, Spooren and Noordman (1992) distinguish between *relational* coherence, denoting the semantic or logical relations between discourse segments, and *referential* coherence, having to do with repeated reference to the same entities or events. Relational coherence can be cued by connectives; anaphora is the primary device for referential coherence. Inferencing is a third major mechanism in creating coherence.

Research on the development of discourse integration has been relatively sparse. Relational coherence, in particular, seems to be underexposed (see Evers-Vermeul, 2000). Work on referential coherence, notably the use of pronominal anaphora, has a longer history; a classic example is Karmiloff-Smith (1980). In this paper we focus on another type of anaphora, viz. nominal ellipsis. We want to determine whether children understand ellipsis, particularly, whether they know that a proper understanding of ellipsis depends on discourse linking.

Discourse integration has been argued to be a late achievement in language acquisition. Karmiloff-Smith (1980) found that children under age 6 largely rely on deictic resolution of definite noun phrases and pronouns in a story context, rather than anaphoric interpretation. Avrutin and Coopmans (2000) studied 4-year-olds’ ability to make bridging inferences (Clark & Haviland, 1977), specifically in contexts where a definite DP can refer back to a DP containing a different noun, in a preceding sentence. They found that children often fail to do this, which they ascribe to a lack of computational resources to complete necessary operations at the level of the discourse model (but see Van der Weert, 2002).

Krämer (2000) studied Dutch-speaking children’s interpretation of indefinite object DPs in base (low) and scrambled (high) positions. Her data show that most children interpret low object DPs in an adult-like manner, but that they treat scrambled indefinite objects in the same way as non-scrambled ones. Krämer claims that it takes until age 7 or 8 until children begin to perform adult-like with regard to scrambled indefinites. She argues that a high indefinite DP needs to be linked to an element in the preceding discourse context (this is not so for low indefinites). Linking depends on the ability to make bridging inferences. To account for young children’s failure to correctly link high indefinite DPs, she assumes that they are not yet capable of discourse integration through bridging.

These results suggest that young children generally have a poor understanding of discourse (including dialogue); they cannot integrate. This strikes us as implausible – it seems strange to deny that a child can integrate the semantics of discourse as soon as a story holds a child’s attention. Two other versions of the ‘no integration’ hypothesis could be developed. First, one could plausibly argue that sentence by sentence the child creates a semantic or conceptual representation, but forgets the syntax. Consequently, the child would be able to connect meanings, but linking on the basis of grammatical devices, such as pronouns or determiners, would be impossible. Alternatively, one could argue that visual context provides support that makes grammatical connections unnecessary (Avrutin & Coopmans, 2000). When presented with a sequence like (1), a child who is incapable of storing the syntactic form of the first sentence cannot connect ‘he’ to ‘John’. But if there is a picture available, ‘he’ could be connected directly to the picture and no discourse information would be necessary, making it seem that the child has integrated the discourse when in fact it has not occurred.

(1) John has a new coat. He put the coat on.

The present study was set up to demonstrate that, counter to claims made in the literature, young children can link sentences into a discourse, and that they make use of syntactic devices to do so. We focus on nominal ellipsis, a type of discourse anaphora that depends crucially on syntactic representation and integration. On the assumption that the principles determining syntactic structure building are available as early as age 2, we predict that young children can interpret this type of ellipsis through syntactic integration, i.e., the identification of syntactically structured information in the immediately preceding discourse. It follows that failures of discourse integration must be traceable not to an absence of the syntactic integration mechanism, but to the child choosing an incorrect grammatical representation where Universal Grammar allows variation (Abdulkarim & Roeper, 1997; Guo, Foley, Chien, Chiang, & Lust, 1997; Matsuo & Duffield, 2001; Thornton & Wexler, 1999).

Ellipsis is ubiquitous in natural language. Hardly a single dialogue fails to use it, and it is a frequent phenomenon in child-directed speech, as a perusal of the CHILDES database demonstrates. Children must cope with it, and, according to our impression, usually do so without many problems, cf. the dialogue excerpt from the Sarah corpus (Brown, 1973; MacWhinney, 2000) in (2).

(2) *CHI: I drink it all up .
*CHI: give me some more .
*CHI: a lot .
*MOT: I don't see any more .
*CHI: yes you do .
*MOT: want a little milk ?
*MOT: want some ?
*CHI: (a)n(d) shake it all up .
*CHI: a bigger one ?
*MOT: mmhm .

Note that the expressions ‘some’, ‘any’, ‘all’, ‘a lot’, ‘one’, ‘more’, ‘you do’, and even ‘mmhm’ all call for reconstruction and it does not seem to be a problem for either the child or the adult. It is possible, however, that in this and similar dialogues, the non-linguistic context helps the child to connect the sentences into a coherent representation, as pointed out above. Clearly, showing that children can do discourse integration on the basis of linguistic information calls for a properly controlled experiment.

Our study uses a paradigm pioneered by Lyn Frazier and her colleagues (Frazier et al., 2003). Adapting a setup used by Wijnen and Kaan (2003) in an online experiment with adults, we had our target sentences contain a nominal ellipsis marked by a bare cardinal, as in example (3).

(3) Five ships appeared on the horizon. Three sank.

Frazier et al. and Wijnen and Kaan have shown that adult perceivers prefer a ‘presuppositional’ or ‘forward-directional’ (Hendriks & de Hoop, 2001) reading of the bare cardinal. This means that they take the elided nominal element to refer to the set defined in the introductory sentence (i.e., three of the five ships that appeared on the horizon). If the second sentence is manipulated in a way that blocks the preferred meaning (for instance by changing the cardinal into one that exceeds the size of the context set), processing difficulty arises.

In the present study, short discourses analogous to the one in (3) were presented together with a picture. The critical sentence was turned into a question, which the participant was encouraged to answer (see ex. 4). Crucially, the combination of picture and story was set up in a manner as to be able to differentiate between responses based on visual support and those based on discourse integration.

(4) Here’s a playground.
It’s great to do all kinds of funny things when you’re out in the playground,
like swinging, making a sand castle or climbing on the monkey bars.
There are some kids playing in the sand box.
Are two upside down?

In one condition – the control condition – the discourse in (4) was combined with a picture of a playground containing the usual paraphernalia, including a sandbox. Crucially, the picture displayed two children in the sandbox who are standing on their heads. In another condition, we presented the same discourse with a picture identical to the first in all relevant respects, except that the two children who are upside down are *outside of the sandbox*. The third condition also involved exactly the same discourse, but this time combined with a picture in which two *adult* figures outside of the sandbox are upside down. Note that in each condition the picture displays two figures that are upside down. This implies that if the child interprets the bare cardinal deictically, taking it to mean something like ‘two anything’, she will give a positive response in all three conditions. However, if the child knows that the bare cardinal marks an elided nominal element that needs to be reconstructed on the basis of the previous discourse, she is predicted to say ‘yes’ in the first (control) condition only, and ‘no’ in the other two.

The difference between the latter two conditions (children upside down outside the sandbox vs. adults upside down) allows us to explore the degree to which reconstruction is affected by the syntactic status of the antecedent. In the second condition described above, a correct ‘no’ answer can only be supplied if the child reconstructs both ‘kids’ and ‘in the sandbox’. We call the phrase ‘kids’ the *argument* because the sentence is uninterpretable without some noun, and ‘in the sandbox’ an *adjunct* because it is syntactically and semantically optional. In the third condition, reconstruction of the argument only (kids) suffices to give a ‘no’ answer.

2. Experiment I: English

2.1 Method

Participants. The experiment involved 28 children recruited from two daycare centers in Amherst, Massachusetts, USA. Mean age of the children was 53.6 months (= 4;6 approximately; range: 40-69 months)

Materials. We used 15 stories analogous to the one in example (4), each combined with one of three different pictures, to create the three experimental conditions as explained above, which we label ‘control’ (C), ‘argument & adjunct reconstruction’ (A) and ‘argument only reconstruction’ (E), respectively. The experimental items were distributed across three lists so that each list contained only one version of each item. The experimental items were interspersed with 10 filler items, a number of which were explicitly designed to elicit ‘no’ answers.

Procedure. We used a variant of the truth value judgment task (Crain & Thornton, 1998). Children were presented picture-story combinations in a setting that had nearly all of the characteristics of normal picture book reading. Children were encouraged to answer the critical question (and other questions) in any way they saw fit. All participants were tested individually by one experimenter in a quiet area in the day care center. A second experimenter was present to make notes and manipulate a video camera, which was positioned unobtrusively in a corner. Just before the actual experiment, the experimenter would informally explore whether the child understood the cardinals (two, three and four) used in the experimental items. It turned out that all children did.

2.2 Results and discussion

The overviews of the results and the analyses refer to averaged percentages of ‘yes’ conditions per experimental condition. Four children exclusively supplied ‘yes’ responses, irrespective of condition. We felt that these children did not understand what was asked of them, and therefore excluded their data from the analyses.

<i>A: Arg + Adj reconstr</i>	<i>C: Control</i>	<i>E: Arg reconstr</i>
35.6%	84.4%	26.7%

Table 1: Average percentages of ‘yes’ responses across conditions.

Table 1 displays the average percentages of ‘yes’ responses as a function of condition. A one-way repeated measures ANOVA yields a significant effect

($F(2,16) = 37.21, p < .001$), indicating that the proportion of yes-responses is reliably lower in conditions A (argument + adjunct reconstruction) and E (argument only reconstruction) than in the control condition C. Recall that, given the visual information in the picture, a deictic strategy would yield a majority of ‘yes’ responses in each of the conditions. Strikingly, the pattern of responses comes close to what we expect for adults. This suggests that the children (average age 4;6) know that the nominal ellipsis must find an antecedent in the preceding discourse, and are furthermore able to reconstruct that antecedent.

The percentage of ‘yes’ responses is somewhat lower in the argument reconstruction condition than in condition A (argument + adjunct reconstruction). This might suggest that performance is better (i.e., more adult-like) in condition E than in condition A. However, an ANOVA indicates that the difference is not reliable ($F(1,17) = 2.47, p = .134$). Nonetheless, we see something of a trend here, which is underscored by a look at individual children’s response patterns. In fourteen of the participants a difference in average score between conditions A and E was observed, and in 10 of these cases, performance on condition E was more adult-like (lower percentage of ‘yes’) than on condition A.

	<i>A: Arg + adj rec</i>	<i>C: control</i>	<i>E: Arg reconstr</i>
<i>42-56 months (n=8)</i>	40%	84%	36%
<i>57-67 months (n=10)</i>	30%	85%	15%

Table 2: Average percentage of ‘yes’ responses broken down over condition and age group.

In order to see whether there is a developmental trend, we divided the subjects in a younger group (42-56 months) and an older group (57-67 months). Table 2 shows the corresponding breakdown of the data. A two-way repeated measures ANOVA yielded a significant main effect of condition, parallel to the one found in the undivided data ($F(2,15) = 41.77, p < .001$). The effect of age group, however, was not significant ($F(1,16) = 2.29, p = .15$), nor was the interaction ($F < 1$). Still, a visual inspection of the results in Table 2 suggest a trend: The percentage of ‘yes’ responses in conditions A and E is higher in the younger group than in the older group, and the difference appears to be largest in condition E. The latter observation is suggestive of a difference in degree of difficulty between the conditions A and E (in line with the individual data presented above), but obviously, we cannot draw firm conclusions in the absence of statistical support. Possibly, the absence of a significant effect is due to a lack of power.

Summarizing, the results of this study – i.e., the significant difference in percentages of ‘yes’ responses between on the one hand the control condition (C), and, on the other, the two experimental conditions (A, E) – indicate that four-year-old children (1) recognize nominal ellipsis (as marked by a bare cardinal); (2) know that such ellipsis requires a *discourse* antecedent (not just any referent); and (3) are capable of reconstructing this antecedent. This finding clearly opposes the suggestion that children below age 6 are incapable of discourse integration.

On the basis of this reasoning, we predict that Dutch bare cardinals *with* ‘er’ allow a less stringent interpretation – argument reconstruction only – than English bare cardinals *without* ‘there’. This predicts that Dutch children (and adults, for that matter) will find a positive reply (yes) to the question ‘Are there two upside-down?’ in the context of the discourse in (5), combined with a picture in which two kids *outside the sandbox* are upside down more acceptable than English-speaking children (and adults).

On the other hand, Dutch ‘er’ is notoriously multifunctional, and acceptability judgments on its presence or absence in relevant contexts vary widely over dialects and individual speakers (Grondelaers & Brysbaert, 1996). In an attempt to get a firmer grip on the effect of ‘er’/‘there’, our second experiment included an additional set of materials, in which nominal ellipsis did not need ‘er’ support. This is possible when if the elided nominal element is indicated by an inflected adjective, as exemplified in (9).

- | | | |
|-----|--|--|
| (9) | Hier is Nijntje.
Er staan allemaal dieren bij haar.
Kijk, er zijn hondjes
<i>Heeft de bruine een bot?</i> | ‘here is Nijntje’
‘there are all kinds of animals with her’
‘look, there are doggies’
‘does the brown (one) have a bone?’ |
|-----|--|--|

If our assumption about the function of ‘there’/‘er’ is on the right track, the prediction should be that discourses like (9) combined with a picture that conforms to a *parallel* interpretation of the ellipsis (e.g. one in which a brown cat has a bone) will yield fewer positive responses than those with critical sentences that contain ‘er’.

Our experiment now entails reference to grammar particular representations in reconstruction. Since these distinctions must be acquired, both the English and the Dutch child may not begin with the right assumption. Is the default assumption that children begin with full reconstruction – treating the numeral as an anaphor – or with the contrastive parallelism that *there* invites? We will argue that where deviation occurs, the child chooses less reconstruction, although the principle of reconstruction seems always available.

3.1 Method

Participants. 47 children (27 boys), recruited at day care centers in Utrecht (The Netherlands) and surrounding towns took part in the experiment. Their mean age was 41.5 months (range 28-57), i.e., 3;6 approximately.

Materials. We used two sets of materials. The first consisted of Dutch adaptations/ translations of the English used in Experiment I. All of these contained the adverbial ‘er’ (there). A new set of items was added in which the nominal ellipsis was marked by a bare (inflected) adjective such as ‘de bruine’ (the brown (one)). In such constructions, the grammar of Dutch does not allow ‘er’. For these new items, pictures were made which instantiated the A, C and E conditions, analogously to what happened in the original items with bare cardinals. For instance, for the discourse given in (8) above, the A condition involved a picture in which a

white dog had a bone, and condition E was constructed by presenting a picture in which a brown *cat* had a bone. Note that there is a difference between these operationalizations and those in the original set of materials, in that the picture in the A condition in itself supplies sufficient information to give an adult-like ‘no’ answer (since there is no brown entity with a bone in the picture).

The old and new items were intermixed and distributed over three lists according to the counterbalancing procedure described for Experiment I. They were interspersed with 9 filler items adapted from the English originals.

Procedure. The experimental procedure was identical to what was done in the English experiment, except that no second experimenter was present.

3.2 Results and discussion

Eight of the children (4m, 4f) virtually only responded with ‘yes’ to virtually all questions, irrespective of condition. This was judged to be the result of a failure to understand the task, and therefore their data were not included in the analysis.

Table 3 gives the average percentages of ‘yes’ responses, broken down over condition and item set. Just as in Experiment I, the effect of condition is significant in a one-way repeated measures ANOVA ($F(2,33) = 65.63, p < .001$), suggesting, that these Dutch children, just like their English-speaking peers, recognize ellipsis and know that it should be interpreted through the reconstruction of elements in the preceding discourse. The effect of item set was marginally significant ($F(1,34) = 4.02, p = .053$), supporting the impression that the percentage of ‘yes’ responses was slightly lower in the new (no ‘er’) items (51%) than in the old ones (59%), overall. This result is in line with our suggestion that ‘er’ facilitates a parallel interpretation of the elided element. It should be noted, though, that the effect occurs irrespective of condition, whereas we expected it to surface in A and E conditions. There was no reliable interaction between condition and item set ($F < 1$), however.

	<i>A: Arg+adj reconstr</i>	<i>C: Control</i>	<i>E: Arg reconstr</i>
<i>old items (er)</i>	45.3%	93%	38.7%
<i>new items (¬ er)</i>	31%	88.1%	33.8%
<i>all items</i>	38.1%	90.6%	36.3%

Table 3: Average percentages of ‘yes’ responses broken down over condition and item set.

The results of Experiment I suggested that argument-only reconstruction is less difficult than argument-plus-adjunct reconstruction – we observed slightly more non-adult-like ‘yes’ responses in condition A than in condition E. The same trend is observed in the present data set, but only for the ‘er’ items. Recall, however, that condition A differs on a crucial point between ‘er’ and non-‘er’ items. In the new items, the visual information in condition A is by itself sufficient to base a ‘no’ answer on, whereas in the old items, inspection of the picture alone (ignoring the preceding discourse) would suggest a ‘yes’. In that sense, the new condition A was easier than the old one. Obviously, the putative effect of the presence or absence of ‘er’ is confounded by this difference.

younger children	<i>A: Arg+Adj reconst</i>	<i>C: Control</i>	<i>E: Arg reconst</i>
<i>old items (er)</i>	47.2%	93.6%	42.6%
<i>new items (¬ er)</i>	38.5%	84.6%	52.6%

older children	<i>A: Arg+Adj reconst</i>	<i>C: Control</i>	<i>E: Arg reconst</i>
<i>old items (er)</i>	44.2%	92.7%	36.4%
<i>new items (¬ er)</i>	26.5%	90.2%	22.7%

Table 4: Average percentages of ‘yes’ responses, broken down over condition, item set and age group. Younger children: 32-38 months (mean 35.9, n = ??); older children: 39-57 months (mean 45.7, n = ??).

To explore whether a developmental trend is present, we divided the subjects over two age groups, and ran a three-way (condition * item set * age group) repeated measures ANOVA. A summary of the results is given in Table 4. The main effect of age group was not significant ($F(1,33) = 2.79, p = .104$), nor did age group interact significantly with any of the other factors. Nonetheless, a visual inspection of Table 4 leaves the impression that the percentage of ‘yes’ responses in the critical conditions A and E decreases with increasing age, particularly in the new items.

4. Cross-linguistic comparison

In order to evaluate whether language makes a difference, we entered the data of the Dutch and American children in one two-way (condition * language) ANOVA. For the Dutch children, we only included responses to the items that have counterparts in the original English set – i.e., the bare cardinal items. Thus, we compare Dutch items with ‘er’ with English ones without ‘there’. The percentages of ‘yes’ responses are given in Table 5. The overall main effect of condition turned out to be significant ($F(2,53) = 61.49, p < .001$), which basically repeats what we have seen before. The main effect of language is significant as well ($F(1, 54) = 4.54, p < .05$), reflecting the overall higher amount of ‘yes’ responses in the Dutch version of the experiment. The interaction was not significant ($F < 1$).

Note that the higher proportion of ‘yes’ responses in the Dutch version of the experiment is an overall effect, present even in the control condition, which would seem to be at odds with a direct effect of language. For some reason, the Dutch children simply seem more inclined to say ‘yes’ than the American children, and this may well have to do with the age difference; the US children are on average nearly one year older than the Dutch children.

	<i>A: Arg+Adj rec</i>	<i>C: control</i>	<i>E: Arg reconst</i>
<i>Dutch</i>	45.3%	93.0%	38.7%
<i>English</i>	35.6%	84.4%	26.7%

Table 5: Average percentages of ‘yes’ responses broken down over condition and language.

The younger group of American children roughly falls within the same age range as the older Dutch children (42-56 months and 39-57 months, respectively). To test whether age is the factor that produced the overall difference between American and Dutch children summarized in Table 5, we ran a new ANOVA, after excluding the older American and younger Dutch children. The results are given in Table 6. As before, the main effect of condition is significant ($F(2,30) = 29.02, p < .001$), but the main effect of language, witnessed in the previous cross-linguistic analysis, has disappeared ($F < 1$). Neither is there an interaction of condition and language.

	<i>A: Arg+Adj rec</i>	<i>C: control</i>	<i>E: Arg reconstr</i>
<i>Dutch: 39-57 mo</i>	44.2%	92.7%	36.4%
<i>US: 42-56 mo</i>	40%	84%	36%

Table 6: Average percentages of ‘yes’ responses broken down over condition and language, comparing younger American and older Dutch children.

5. General discussion

The evidence collected in our experiments clearly shows that young children – as young as 3 years of age – know how to interpret nominal ellipsis. Given that ellipsis is ubiquitous in natural dialogue and that children generally do not seem to be puzzled by it, this may be unsurprising, but in fact it is an amazing achievement. Presented with a picture of a playground and a discourse like ‘*Some kids are playing in the sandbox. Are two upside down?*’, the child realizes that ‘two’ cannot refer to any two entities that happen to be upside down in the picture. Rather, she knows that the bare cardinal is an anaphor, and that a certain cardinality (two-ness) is predicated over an unexpressed element, which is to be reconstructed from what the adult experimenter has *said*, not from what she *shows*. In other words, the child demonstrates an ability to integrate discourse. We believe that children are capable of this remarkable feat because their syntactic competence helps them. Our claim is that children know that elements like bare cardinals mark incompletely specified syntactic structure, which needs to be filled in with material from the immediately preceding discourse; in other words, they possess a mechanism for syntactic integration.

Our experiments contained a condition in which reconstructing an argument only was sufficient to see that the experimenter’s question should be answered in the negative, and one in which both the argument and an adjunct were crucial. Assuming that the child’s answer is based on syntactic reconstruction, as we do, it would seem likely that the latter case (condition A) is more difficult – and yields more errors – than the former. It turns out that, although we think we see the beginnings of a trend in this direction, we cannot decide in favor of this prediction. Further, more careful experimentation is needed.

We expected an effect of variation in grammar on reconstruction in ellipsis. Specifically, we speculated that there-insertion would promote parallel interpretation of the ellipsis, rather than a presuppositional interpretation (which means that with ‘there’ present, perceivers would tend to give an affirmative response to the question in (4) when the picture shows two children upside down outside of the sandbox).

The slight difference between the English (without ‘there’) and Dutch (‘er’ present) results suggests that children are sensitive to the properties ‘there’. It should be kept in mind, though, that neutralizing the age effect quite effectively suppressed the difference between Dutch and American children.

In the Dutch items without ‘er’, parallel interpretation was possible in condition E (where e.g. a brown *cat* is shown to have a bone). The number of ‘yes’ responses in this condition is lower than that in the old items (with ‘er’) in condition A. Although this observation is suggestive, we are reluctant to put very much weight on it, largely because the tendencies cannot be backed up statistically, but also because the old and new items differ in too many respects to make a direct comparison possible. At this point, we can only conclude that the results seem to hold some promise, warranting new and more carefully controlled experiments.

Such experiments should be guided by the following considerations. The child must acquire the distinction between the various functions of ‘there’: the locative, the presentational and the existential. In Dutch, the presentational-existential distinction is not marked, so ‘er’ remains ambiguous. In English, very subtle evidence comes into play, as the contrast in (10) reveals.

- (10) a. presentational: There’s no unicorns. (not here/now)
(as in: no unicorns in this painting)
- b. existential: There are no unicorns. (don’t exist)

Our prediction, given that the principle of syntactic discourse integration is present from the outset, is that children will immediately integrate new syntactic knowledge into the application of discourse integration. Therefore, when the child acquires the distinction between presentational and existential, it would appear instantly under the parallelism requirement. The general research program which emerges is to explore the claim that every new syntactic distinction will be immediately available to discourse reconstruction.

In conclusion, our study demonstrates that young children are capable of discourse integration, in contradiction to claims made earlier. The contrast between the findings reported here and previous results may be related to a difference in linguistic device. We studied underspecified, non-overt anaphoric devices, rather than pronouns or determiners which figure in the literature cited. On the other hand, even in dealing with pronouns, children may be more advanced than has been generally thought. Moore, Roeper, Asplin and Wagner (work in progress) presented children with simple visual scenes in which for instance Bert plays basketball and Ernie baseball. They provide the following text: ‘*Bert is playing basketball. Is he playing baseball?*’. If the child were simply using visual context, the answer should be ‘yes’, since there is a visible ‘he’ playing baseball. But in fact the children say ‘no’, indicating that they link ‘he’ to Bert. It would seem that some of the earlier studies have seriously underestimated children’s early competence.

References

- Abdulkarim, L., & Roeper, T. (1997). Economy of representation: Ellipsis and NP reconstructions. In R. Shillcock (Ed.), *Language Acquisition: Knowledge representation and processing*. Edinburgh, UK: Human Communication Research Center.
- Avrutin, S., & Coopmans, P. (2000). Children who build bridges, *Proceedings of BUCLD* (Vol. 24, pp. 80-91). Somerville MA: Cascadilla Press.
- Brown, R. (1973). *A first language: the early stages*. Cambridge, MA: Harvard University Press.
- Clark, H., & Haviland, S. (1977). Comprehension and the given-new contract. In R. Freedle (Ed.), *Discourse Production and Comprehension* (pp. 1-40). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Crain, S., & Thornton, R. (1998). *Investigations into universal grammar: a guide to experiments on the acquisition of syntax and semantics*. Cambridge MA: MIT Press.
- Evers-Vermeul, J. (2000). De complexiteit van connectiefverwerving [The complexity of connective acquisition]. *Nederlandse Taalkunde*, 5(3), 250-271.
- Frazier, L., Clifton, C., Rayner, K., Deevy, P., Koh, S., & Bader, M. (2003). *Interface problems: Structural constraints on interpretation?* Unpublished manuscript, Amherst MA.
- Grondelaers, S., & Brysbaert, M. (1996). De distributie van het presentatieve er buiten de eerste zinsplaats. *Nederlandse Taalkunde*, 1, 280-305.
- Guo, F., Foley, C., Chien, Y.-C., Chiang, C.-P., & Lust, B. A. (1997). A cross-linguistic study of Chinese and English children's first language acquisition of VP Ellipsis structures. *Papers on Language Acquisition: Cornell Working Papers in Linguistics*, 15, 20-38.
- Hale, K., & Keyser, J. (2000). *There-insertion unaccusatives*. Unpublished manuscript, Cambridge, MA.
- Hendriks, P., & de Hoop, H. (2001). Optimality theoretic semantics. *Linguistics and Philosophy*, 24, 1-32.
- Karmiloff-Smith, A. (1980). Psychological processes underlying pronominalization and non-pronominalization in children's connected discourse. In E. Ojeda (Ed.), *Papers from the parasession on pronouns and anaphora*. (pp. 222-250). Chicago IL: Chicago Linguistics Society.
- Krämer, I. M. (2000). *Interpreting indefinites : an experimental study of children's language comprehension*. Unpublished PhD, Utrecht University, Utrecht.
- MacWhinney, B. (2000). *The CHILDES project: tools for analyzing talk* (3rd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.
- Matsuo, A., & Duffield, N. (2001). VP-ellipsis and anaphora in first language acquisition. *Language Acquisition*, 9(4), 301-327.
- Sanders, T. J. M., Spooren, W. P. M., & Noordman, L. G. M. (1992). Coherence relations in a cognitive theory of discourse representation. *Cognitive Linguistics*, 4, 93-133.
- Thornton, R., & Wexler, K. (1999). *Principle B, VP Ellipsis, and Interpretation in Child Grammar*. Cambridge, MA: MIT Press.
- Van der Weert, C. F. (2002). Nature or nurture: structured information knowledge. *Reading Working Papers in Linguistics*, 6, 119-132.
- Wijnen, F., & Kaan, E. (2003). *Interpreting anaphoric quantifiers in context*. Unpublished manuscript, Utrecht, NL.