

Children's Understanding of Distributivity and Adjectives of Comparison

Anna de Koster, Jakub Dotlačil, and Jennifer Spenader

1. Introduction

A fundamental property of human language is its ability to express information about quantities of entities, rather than only about individuals. This can be achieved in various ways, such as using quantificational expressions (each/every boy), numerical expressions (three boys) or plurals (the boys). Take for example the following sentence: *The boys are pushing a car.* This sentence allows multiple interpretations: Are the subjects acting collectively as a group (the collective interpretation, Figure 1) or separately as individuals (the distributive interpretation, Figure 2)?



Figure 1. Collective Interpretation



Figure 2. Distributive Interpretation

The acquisition of distributivity has been studied for decades, but many questions are still unanswered. This project addresses two of them. First existing research has often focussed on the acquisition of overtly distributive quantificational determiners like *Each* and *Every*, but less work has investigated preferences and dispreferences for distributive interpretations with other determiners. Study 1 examines to what degree improvements in the interpretation of *Each/Every* correlate with more adult-like preferences for collective interpretations with the definite article *The*. Second, distributive

*Anna de Koster, University of Groningen, a.m.b.de.koster@rug.nl, Jakub Dotlačil, University of Groningen, j.dotlacil@gmail.com, Jennifer Spenader, University of Groningen, j.k.spenader@rug.nl. The first author is supported by a grant from the Netherlands Organization for Scientific Research. We would like to thank all the participants and in special OBS De Sterrensteen.

interpretations are required for the correct interpretation of other linguistic expressions, but little research has examined children's intuitions with these constructions. Study 2 addresses this issue by investigating children's preferences in interpreting the adjective of comparison *Different*, which has a distributivity component. Because the same children were tested in both studies, we can compare how interpretations with *Different* relate to distributivity intuitions with *Each/Every* and *The*.

2. Background

Consider the following two sentences:

- (1) **Every/each** boy is pushing a car.
- (2) **The** boys are pushing a car

Adults interpret sentences like (1) only with a distributive interpretation. On the other hand, adults strongly prefer collective interpretations for sentences like (2) where the subject is a definite plural (Frazier et al, 1993; Kaup et al, 2002). Children instead tend to prefer distributive interpretations regardless of noun phrase type (Syrett and Musolino, 2013). How do children then develop adult preferences?

Dotlačil (2010) hypothesized that definite plurals conversationally implicate their collective meanings. In short, adults have both distributive and collective interpretations available for definite plurals, but reason that if the speaker had meant a distributive interpretation she would have used an explicit distributivity marker like *Every* or *Each*. This explanation treats distributivity as part of the lexical meaning of *Each/Every*, but treats the collective preferences with plural subjects as derived via pragmatics. Young children (up to the age of 8) are not able to reason about alternative expressions, because they have an incomplete semantic representation of the quantifier *Each/Every* lacking the information that it has distributive meaning. *Each/Every* and *The* mean exactly the same to them, so they can't use the former to exclude distributive interpretations of the latter.

This hypothesis leads to the following prediction: the rate of rejection of *Each/Every* in the collective context will correlate positively with the rate of rejection of *The* in the distributive context. This means that when children learn to understand *Each/Every* they will also start rejecting *The* in the distributive context. Pagliarini et al. (2012) found such a correlation in Italian children for the definite *i/le* and the quantifier *Ciascuno* 'every/each'. In Study 1 we tested this same hypothesis with Dutch participants.

3. Study 1 - Comparing distributivity preferences for *The* and *Each* for children and adults

Study 1 used a truth value judgment task with two factors: Picture (*Distributive, Collective*) and Determiner (*Definite Plural, Each*), using the Dutch definite plural *De* 'The' and the quantifier *Iedere* 'Each'.

3.1. Participants

114 children, divided in five different age-groups from 5 to 9 years old, participated. The children were tested in a quiet classroom. 40 adults who were mainly university students, served as a control group. They performed the experiment online, without the experimenter being present.

3.2. Design and Procedure

We used two types of pictures. They depicted either a distributive context (Figure 3) or a collective context (Figure 4).

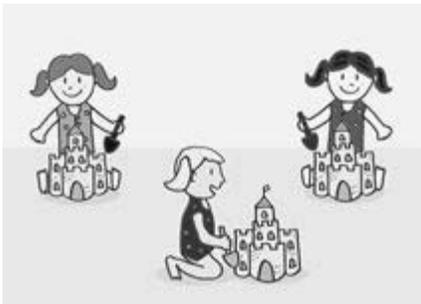


Figure 3. Distributive



Figure 4. Collective

Sentences were of the form Subject-Verb-Indefinite Object, beginning either with the definite plural *De* (*The*) or the quantifier *Iedere* (*Each/Every*):

- (3) **De** meisjes bouwen een zandkasteel.
The girls are building a sandcastle.
- (4) **Ieder** meisje bouwt een zandkasteel.
Each girl is building a sandcastle.

Six verbs were used: *vasthouden*, *dragen*, *duwen*, *trekken*, *wassen* and *bouwen* (in English: 'hold', 'carry', 'push', 'pull', 'wash' and 'build'). The subjects were girls, boys, monkeys or dogs and every item contained a different object.

The 2x2 design thus has four conditions (Table 1). Participants saw six items per condition, plus twelve control items. Items were distributed over four lists and were presented randomly to the participants. They were presented with one picture at a time, while a recorded sentence was played, and asked to verify whether the sentence matched the picture.

Table 1. The different conditions plus adult predictions

Condition	Adult Prediction
Each - Distributive	Yes
The - Collective	Yes
Each - Collective	No
The - Distributive	No

3.3. Results

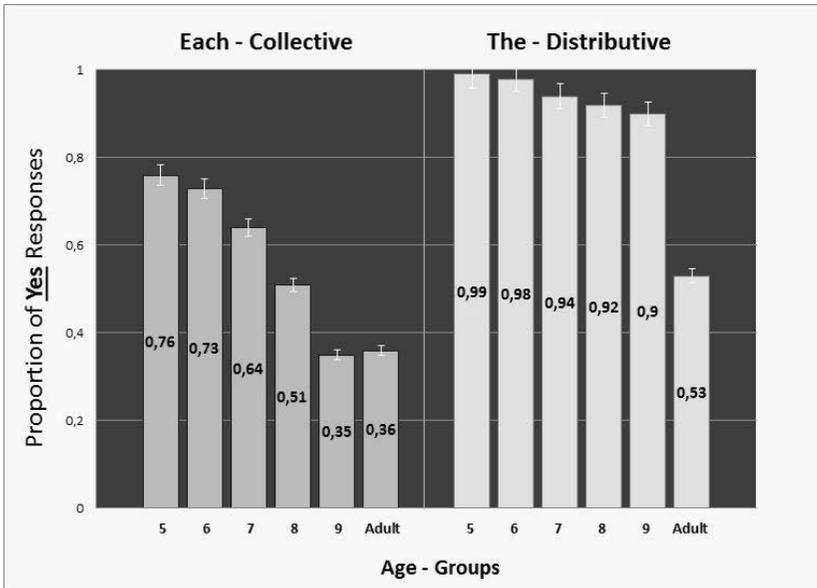


Figure 5. Mean proportion of 'Yes' responses for the *Each-Collective* and the *The-Distributive* condition in Study 1 with standard error bars.

Both the adults and the children consistently accepted *Each-Distributive* and *The-Collective* around 99% of the time even from the age of 5 years old, so we omit the figure for these conditions. Conditions *Each-Collective* and *The-Distributive* showed larger differences between the children and adults, see Figure 5. The results of the *Each-Collective* condition suggest that 5 and 6-year-old children accept the collective interpretation of the quantifier *Iedere* 'Each'. However, starting at age 7, children start rejecting this interpretation,

gradually moving toward the adult interpretation, which is reached at the age of 9. The results of the condition *The-Distributive* show that children from the age of 5 accept the distributive interpretation of the plural definite *De 'The'*. They start rejecting this condition around the age of 8/9 years old. However they are far from adult-like, in contrast to their results for the *Each-Collective* condition. This finding suggests that children learn to restrict *Each* to distributive meanings before they start rejecting *The* in the distributive context.

The results were analyzed using mixed-effect logistic models. The dependent variable was the response and the predictors were: *CONDITION* (with the reference level: *Each-Distributive*), *AGE*, and the interaction of the two. We also included two random effects for the intercept: *PARTICIPANTS* and *VERB*, and one random effect for the slope of the verbs: *CONDITION*. In the final model, *CONDITION Each-Collective* turned out to be a significant predictor ($\beta = -6.33$, $z=-5.8$, $p<0.001$). Condition *Each-Collective* is rejected significantly more than the other conditions. Furthermore, two interactions were significant: *AGE* with *CONDITION Each-Collective* ($\beta = -0.2$, $z=-2.2$, $p<0.01$) and *AGE* with *CONDITION The-Distributive* ($\beta = -0.36$, $z=-3.9$, $p<0.001$). Other conditions and their interactions with age were not significant. These interactions show that as children grow older both Conditions *Each-Collective* and *The-Distributive* are accepted less.

Recall that the implicature analysis for distributivity (Dotlačil, 2010) predicts that children should learn to reject Condition *Each-Collective*, as well as Condition *The-Distributive*. This seems to be confirmed by the descriptive statistics and the results of the mixed effect logistic model. Furthermore, we predict that the correlation between the conditions *Each-Collective* and *The-Distributive* should hold at the level of the individual child.

To examine this prediction, we have to examine the correlation between each child's acceptance of two conditions. We defined 'acceptance' as the number of items that a child accepted in a condition. The correlation between the proportion of items accepted in the Condition *Each-Collective* and the proportion of items accepted in the Condition *The-Distributive* was measured using Spearman's rank correlation. A significant positive correlation was found between the two conditions ($\rho=0.3$, $p<0.001$). None of the other conditions correlated significantly with Condition *Each-Collective*.

Figure 6 shows the correlation between the two conditions, by showing how many times each child accepted Condition *Each-Collective* (x-axis) and Condition *The-Distributive* (y-axis). Note that there are no data points in the upper left corner of the graph. Data points in that corner would represent children that reject *Each* in the collective context, but accept *The* in the distributive context. According to our hypothesis, that combination should be impossible, and the empty upper left corner thus serves as evidence for the prediction that learning *Each* precedes learning to reject *The* in the distributive context.

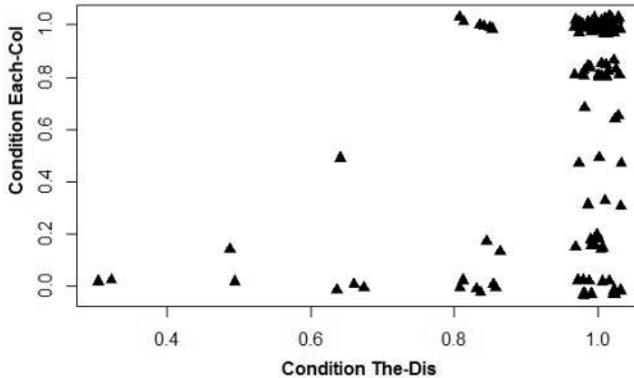


Figure 6. Correlation between the acceptance rates of Conditions Each-Collective and The-Distributive: The triangles represent the individual children.

3.4. Discussion

We found our predicted positive correlation that children start rejecting *Each* in the collective context before they start rejecting *The* in the distributive context, and this holds even at the level of the individual child.

However, the acceptance rate of the *Each-Collective* condition for adults turned out to be higher than expected. We expected adults to overwhelmingly reject this condition but they still accepted it in 36% of the cases. When we compare this with the results of Pagliarini et al. (2012) we see that they found a lower acceptance rate of only 9%. One might think that this difference could be due to the different verbs we used and that some of them might be easier to interpret as collective (e.g. *to build*). This would predict that for the *Each-Collective* condition, the factor *VERB* should be a significant predictor. We tested this in a model but found no effect, suggesting that different verbs played little to no role in this condition.

Another finding that needs some discussion is the unexpectedly high 50% acceptance rate of the adults in condition *The-Distributive*. However, these results are consistent with the results of Pagliarini et al. (2012) and are in line with our hypothesis since it is known that rates of implicature calculation vary across different lexical items and most implicatures are calculated in around 30% to 70% cases (c.f., van Tiel et al., 2016).

3.5. Other interpretations requiring distributivity

Adjectives of comparison (AOCs), e.g. *different*, *same* and *similar*, are used to compare two or more elements and in one of their interpretations, they are

parasitic on distributivity, as discussed below. Study 2 examines the relationship between distributivity and the Dutch adjective of comparison *andere* 'Different' and the use of *Each* or *The*. Consider the following sentence:

(5) The girls pushed a Volvo. Each boy pushed a different car.

(5) is ambiguous. First, the boys might have pushed cars distinct from each other. This is the sentence internal reading, where *Different* is interpreted as signalling distinct cars paired to each boy. There is also a sentence external reading, where all boys push a car distinct from the Volvo. In the sentence external reading the car in the current sentence is compared to an element, here the Volvo, mentioned previously in the discourse, while in the sentence internal reading cars are compared between the same subjects. Carlson (1987) proposes that the one-to-one pairing necessary for the internal reading is a distributive interpretation (see also Moltmann, 1992, Brasoveanu, 2011). Thus internal readings require distributive interpretations.

Carlson's (1987) claim is supported by parallelisms that have been found between the 'strength' of distributivity markers and the acceptability of sentence-internal 'different' licensors. Brasoveanu and Dotlačil (2012) found in a questionnaire study on English that the degree to which different NPs license distributive readings can be organized on a scale and that this scale of distributivity markers is the same as the scale for sentence-internal *Different* licensors.

Acceptable **EACH > ALL > {THE, NONE}** **Unacceptable**

Their results showed that when the subject NP appeared with *Each*, the sentence internal reading of *Different* and a distributive reading were fully accepted. When the subject NP was *The*, the same readings were degraded. These results then support the claim that the sentence-internal reading of *Different* requires distributivity to be licensed. For children, these results suggest that sentences with *Different* with the NP subject *Each* will be treated differently from those with the NP subject *The*. Given the results from Study 1, we would predict that children will first need to restrict the *Each* sentences to sentence internal readings before they will learn to reject internal readings with *The* as the subject NP.

Study 2 examines these predictions by testing children's interpretation preferences for the adjective of comparison *Different* and further compares these results with those of Study 1. We predict that when children show adult interpretations restricting *Each* to distributive meanings, they will also show their mastery of distributivity by preferring the sentence-internal reading of the AOC *Different* in combination with the strong licensor *Each*.

4. Study 2 - The relationship with the adjective of comparison *Different*

Study 2 used a truth value judgment task with two factors: Picture and Sentence. The experiment is conducted in Dutch, with the definite plural *De 'The'* and the quantifier *iedere 'Each'* in both a sentence-internal and a sentence-external context. Additionally we also tested for a correlation between the results of Study 1 and Study 2 to check if the children who showed the adult interpretation of distributivity (rejecting *The* in the distributive context) will also show the adult reading of the adjective of comparison *Different* (rejecting the internal reading of *Different* in combination with sentences with *The*).

4.1. Participants

86 children who participated in Study 1 also participated in Study 2. The children were divided in four different age-groups from 6 to 9 years old. For Study 2 the 5 year olds were excluded because they already had difficulties with Study 1. 52 adults served as a control group. None of the adult participants had taken part in Study 1.

4.2. Design and Procedure

We used two different picture types in this experiment. They were presented as little comics and were accompanied by a recorded context story. The pictures were either 'internal' (Figure 7) or 'external' (Figure 8). A yes answer to the internal pictures corresponds to an internal reading and a yes answer to the external picture corresponds to an external reading, hence the names of the pictures.

Figure 7 shows the internal picture. The actual pictures were in color. The black arrows were not present but are used here to indicate the same car in the third and fourth panel. This comic was accompanied by the following context story:

There are three cars, a green one, a blue one and a red one. The girls are pushing the green car one after another. The first boy is also pushing the green car. The second boy is pushing the blue car and the last boy is pushing the red car.

Figure 8 shows the external picture and the context story belonging to this picture is the following:

There are two cars, a green one and a blue one. The girls are pushing the green car, one after another. The boys are pushing the blue car, one after another.

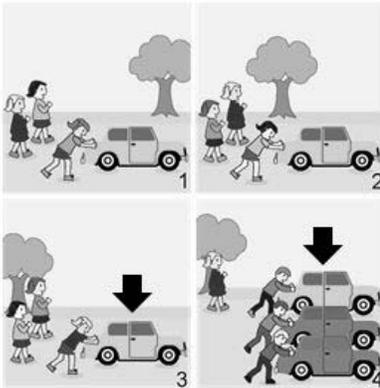


Figure 7. Internal

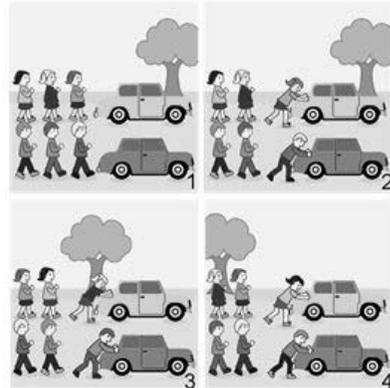


Figure 8. External

The context stories were recorded in Dutch, but for reasons of space we show only the English translations. A target sentence and question followed directly after each context story and started with either *de* 'the' or *iedere* 'each'.

- (6) De meisjes duwden dezelfde groene auto. Klopt het dat **de** jongens een *andere* auto duwden?
*The girls were pushing the same green car. Is it true that **the** boys pushed a different car?*
- (7) De meisjes duwden dezelfde goene auto. Klopt het dat **iedere** jongen een *andere* auto duwde?
*The girls were pushing the same green car. Is it true that **each** boy pushed a different car?*

Table 2. Expected answers for adults for each condition

Sentence / Picture	Internal	External
<i>Each</i>	Condition 1 Yes, because all three boys are pushing a different car.	Condition 3 No, because the boys are all pushing the same blue car.
<i>The</i> Evokes comparison with external element (girls)	Condition 2 No, because one of the boys is pushing the same green car as the girls .	Condition 4 Yes, because the boys are pushing a different car than the girls .

Sentences (6) and (7) were paired with Figures 7 and 8 and exemplify the four different conditions of the 2 x 2 design. Our predictions per condition are shown in Table 2. The same verbs were used as in study 1, we only left out *build* and *wash*, because they were too difficult to depict in the comic style pictures. The subjects and objects remained also the same. Participants were presented with 5 items per condition plus 12 control items. Items were distributed over four lists and were presented randomly to the participants. They were presented with one picture at a time, while a recorded story was played. The target question followed directly after the context story.

4.3. Results

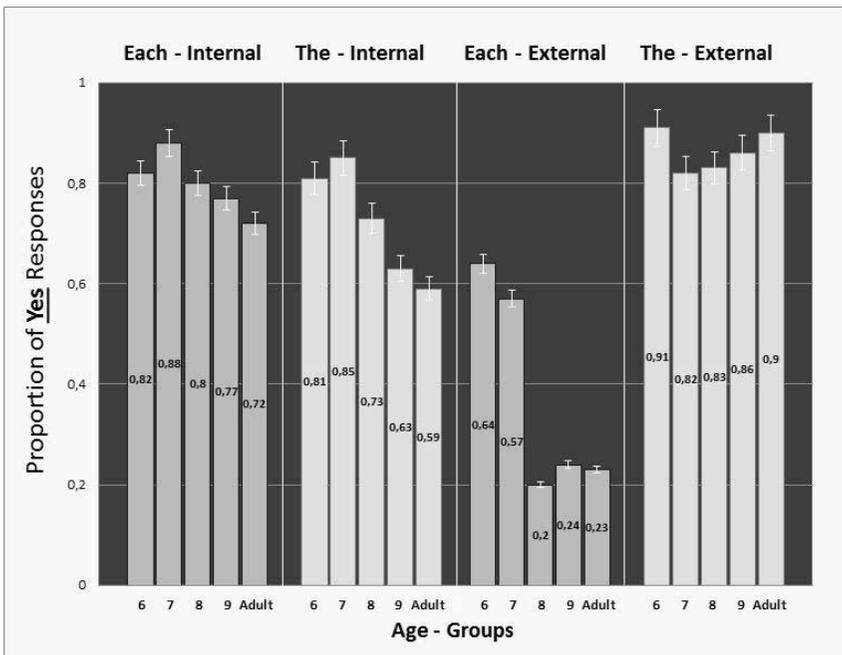


Figure 9. Mean proportion of ‘Yes’ responses for all conditions in Study 2 with standard error bars

Figure 9 reports the mean proportion of ‘Yes’ responses per condition. We start with the internal readings. When comparing the results from the adults in Figure 9 with our predictions shown in Table 2, it becomes clear that adults indeed accept condition *Each-Internal*, but with an acceptance rate of 72% which is slightly lower than expected. We predicted adults would reject condition

The-Internal, because of the expected external reading, but they show an acceptance rate of about 60% which is higher than expected.

In case of the external readings, we see that adults accept *Each-External* only at the rate of 23%, which closely matches our predictions. The 90%-acceptance rate of *The-External* also accords with the predictions. When looking at the children, we see they are adult-like for *Each-External* at the age of 8, and at the age of 6 for *The-External*.

The results were analyzed using mixed-effect logistic models, with the response as the dependent variable. The predictors were: *CONDITION* (reference level: *The-Internal*), *AGE* and the interaction of the two. We also included two random effects for the intercept: *PARTICIPANTS* and *VERB*, and one random effect for the slope of the verbs: *CONDITION*. In the resulting model the condition *Each-External* turned out to be a significant predictor ($\beta = -2.6$, $z=-9.7$, $p<0.001$). The interactions of *AGE* with Conditions *Each-External* and *The-External* were also significant ($\beta = 0.04$, $z=3.0$, $p<0.001$) and ($\beta = 0.09$, $z=4.8$, $p<0.001$). Furthermore, the predictor *AGE* also turned out to be significant ($\beta = -0.07$, $z=-4.8$, $p<0.001$). These results show us that older speakers accept condition *The-External* more than *The-internal*. Condition *Each-External* on the other hand is rejected more. Returning to the question of whether or not there is a relationship between distributivity and the adjective of comparison *Different*, we examined the correlation between each child's acceptance of condition *The-Distributive* from Study 1 and their acceptance of condition *The-Internal* from Study 2. We defined 'acceptance' as the number of items that a child accepted in a condition.

We used Spearman's rank correlation to test for a statistical relationship but found no correlation between the two conditions ($\rho=0.1$, $p>0.1$). The lack of correlation is also visually confirmed in Figure 10 which shows how many times each child accepted condition *The-Distributive* (x-axis) and condition *The-Internal* (y-axis).

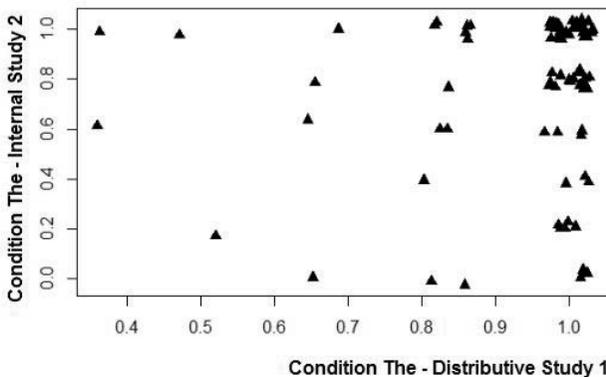


Figure 10. Correlation between the results of Study 1 and Study 2.

4.4. Discussion

The results of Study 2 followed our predictions for external readings. Adults accept external readings with *The* and reject external readings with *Each*. Children are adult-like for *The*-External and show the adult-like interpretation of *Each*-External at the age of 8. Adult responses are in line with our predictions, see Table 2. The responses of children also follow our predictions relatively closely. First, given that children fully accept *The* in its collective interpretation already at the age of 5, we expect *The*-external to be acceptable for young children, which is correct. Second, we saw that children are adult-like in their rejection of *Each*-Collective at the age of 9 and we predict the same pattern for *Each*-External. In this study, we see that they are adult-like in rejecting *Each*-External at the age of 8.

The results for the internal picture on the other hand turned out to be different than predicted. The adults accepted both conditions, but we expected them to reject *The*-Internal items. They showed an acceptance rate of 59%, which indicates an internal reading. However, the 59% might be explained by the fact that this condition is again the condition in which the calculation of the implicature takes place, just like the *The*-Distributive condition from Study 1, which showed an adult acceptance rate of 50%. Rates close to 50% are not unexpected for implicatures.

Another possible explanation for the unexpectedly high rate of acceptance of internal readings with *The* might be due to the pictures. The external pictures were easier to understand and clearer. In the external picture (Figure 8) there is a distinct separation between the boys and the girls. The girls are pushing one car and the boys are pushing one (other) car. The four subpictures are very similar and the order of the events is transparent. These distinctions make a comparison between the boys and girls very salient. This is contrary to the internal picture (Figure 7), in which the comic consists of more difficult subpictures. The separation between the girls and boys is less apparent and the events are harder to understand. It might be the case that the participants just ignored or did not see that the first boy was pushing the same car as the girls.

Also, the target question was asked after the story was played, so the participants might be just using the last picture in their verification process and this last picture shows that the three boys are pushing three different cars.

Thus, a point for future research seems to be to rerun this experiment with different pictures, not in comic style, but just one picture in which the four subpictures of the comics are combined.

5. General Discussion and Conclusion

We will first discuss the positive findings. In Study 1, we confirmed the hypothesis of Dotlačil (2010) and the results of Pagliarini et al. (2012) for Dutch. Rejections of *The* with distributive interpretations arise as the interplay of semantics and pragmatics. We also confirmed previous findings (e.g., Brooks and Braine, 1996) showing that children often interpret *each* collectively. In Study 2, external readings closely confirmed our predictions both for adults and children. In particular, we saw a close match between the acquisition of *Each* and *The* in a collective interpretation and *Each-External* and *The-External*. The only surprising factor is that the adult-like rejection of *Each-External* appears one year earlier than the rejection of *Each-Collective* (8 years vs. 9 years).

Rates of internal readings in Study 2 did not fit well with our predictions. First, it was surprising to see that *Each-Internal* was only accepted at 72% by adults. Notice also that *Each-External* was rejected at the rate of 77% by adults. Assuming that the rejection in the second case was due to the fact that adults expected an internal interpretation, both results would closely converge on the finding that the internal reading with *Each* is only accepted in 75% cases. This is surprising given our theoretical assumptions. It is also surprising given previous findings, in particular, that of Brasoveanu and Dotlačil (2012), which showed that *Each-Internal* was fully accepted in English. This discrepancy might be rooted in the fact that *Ieder 'Each'* is not so strongly distributive in Dutch as it is in English. This is supported by the findings of Study 1, in which *Each-Collective* has the unexpectedly high acceptance rate of 36%. The second surprising finding concerned *The-internal*. This condition was accepted at the rate of 59% by adults, which is high. It goes against the previous finding of Dotlačil (2010) that acceptance rates of sentence internal readings were lower than acceptance rates of distributive readings. Dotlačil (2010)'s study was in Dutch, so language cannot explain the differences. Also unexplained is the discovery that sensitivity to the sentence internal reading is acquired earlier than adult-like intuitions on straightforward distributivity. While children in Study 2 reached adult-like norms in all conditions by the age of 9, they are still not adult-like in their interpretation of *The-Distributive* according to Study 1.

What could explain the differences in the acquisition of *The-Distributive* and *The-Internal*? We see two possibilities: either the semantic features of *Different* that are not related directly to distributivity have an effect, or there is greater flexibility in anaphoric interpretation compared to felicity or grammaticality judgments. Let's examine the first proposal in more detail. The semantic constraints of distributivity shared in distributive interpretations and in sentence internal readings both require individuals denoted by the subject DP to individually take part in an event, and that objects acted upon must be individual and linked one-to-one to each member of the subject set. *Different* however adds a further requirement that the objects acted upon must be distinct from one

another. This additional semantic constraint might temper preferences because it will generally (though not so much in our actual experiment) make it less likely that the speaker's intended interpretation is misunderstood. If *Each* in general leads to examining individual boys (or girls) while *The* encourages treating the boys as a group, *Different* encourages comparing cars, and this additional semantic information might make the intended interpretation in a context less likely to be misunderstood compared to simple transitive sentences. This position would also explain why the rejection of *Each-External* is acquired slightly earlier than the rejection of *Each-Collective*. If this is correct, then it is possible that children in fact use sentences with *Different* to learn the right interpretation of *Each* and *The*. From that perspective, studying the acquisition of adjectives of comparison might give us an important insight into how the interpretation of plural expressions is acquired.

References

- Brasoveanu, Adrian (2011). Sentence-internal different as quantifier-internal anaphora. *Linguistics and Philosophy*, 34(1), 93–168.
- Brasoveanu, Adrian and Dotlačil, Jakub (2012). Licensing sentence-internal readings in English. Springer.
- Brooks, Patricia J., and Martin D.S. Braine. 1996. What do children know about the universal quantifiers all and each? *Cognition* 60:235–268.
- Carlson, Greg N. (1987). Same and different: Some consequences for syntax and semantics. *Linguistics and Philosophy*, 10(4):531–565.
- Dotlačil, Jakub (2010). Anaphora and Distributivity. A study of same, different, reciprocals and others, Doctoral Dissertation, Utrecht Institute for Linguistics OTS, LOT Series.
- Frazier, Lyn, Patch, Jeremy and Rayner, Keith (1999). Taking on semantic commitments, II: collective versus distributive readings, *Cognition*, vol. 70, pp. 87–104.
- Kaup, Barbara, Kelter, Stephanie and Habel, Christopher (2002). Representing referents of plural expressions and resolving plural anaphors, *Language and Cognitive Processes*, vol. 17, no. 4, pp. 405–450.
- Moltmann, Friederike. (1992). Reciprocals and *same/different*: Towards a semantic analysis. *Linguistics and Philosophy*, 15(4), 411–462.
- Pagliarini, Elena, Fiorin, Gaetano, and Dotlačil, Jakub (2012). The acquisition of distributivity in pluralities. In Proceedings of the Annual Boston University Conference on Language Development, volume 2, pages 387–399.
- Syrett, Kristen, and Musolino, Julien (2013). Collectivity, distributivity, and the interpretation of plural numerical expressions in child and adult language. *Language Acquisition: A Journal of Developmental Linguistics*, 20, 259–291.
- Van Tiel, Bob, Van Miltenburg, Emiel, Zevakhina, Natalia, and Geurts, Bart (2016). Scalar diversity. *Journal of Semantics*, 33(1), 137–175.

Proceedings of the 41st annual Boston University Conference on Language Development

edited by Maria LaMendola
and Jennifer Scott

Cascadilla Press Somerville, MA 2017

Copyright information

Proceedings of the 41st annual Boston University Conference on Language Development
© 2017 Cascadilla Press. All rights reserved

Copyright notices are located at the bottom of the first page of each paper.
Reprints for course packs can be authorized by Cascadilla Press.

ISSN 1080-692X

ISBN 978-1-57473-076-0 (2 volume set, paperback)

ISBN 978-1-57473-176-7 (2 volume set, library binding)

Ordering information

To order a copy of the proceedings or to place a standing order, contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, sales@cascadilla.com, www.cascadilla.com