Are Universal Quantifier Errors and Errors with *only* Related?

Jennifer Spenader, Anna de Koster, Bart Hollebrandse, and Petra Hendriks

1. Introduction

Children's errors with universal quantifiers like *all, every* and *each* have been extensively studied. However, much less research has looked at how children interpret other quantifiers and whether or not there might be commonalities in the errors that they make.

In this paper we contribute to our understanding of general quantifier acquisition by presenting results from experiments investigating children's errors with the quantificational adverbial *only.* In three experiments we tested *only* and universal quantifiers with the same children, allowing us to compare their results. Together the results suggest that children's errors with *only* and with universal quantifiers are related.

2. *Only* and Universal Quantifers

Only and universal quantifiers like *all* and *each* are syntactically quite different: *only* is an adverb, while universal quantifiers are determiners. But despite these syntactic differences, there are semantic reasons to believe that children might make similar errors with universal quantifiers and *only.*

Most semantic analyses of quantifiers emphasize that they signal relations between sets (Barwise & Cooper, 1981). For example, in examples (1) and (2) the quantifiers *all* and *only* signal a relationship between a set of bunnies and a set of dancers.

(1) All the bunnies are dancing.
(2) Only the bunnies are dancing.

Both universal quantifiers and *only* describe a particular type of set relationship, namely a subset relation between the set denoted by the subject (the set of bunnies)

* Jennifer Spenader, University of Groningen, j.k.spenader@rug.nl, Anna de Koster, University of Groningen, a.m.b.de.koster@rug.nl, Bart Hollebrandse, University of Groningen, b.hollebrandse@rug.nl, Petra Hendriks, University of Groningen, p.hendriks@rug.nl.

Proceedings of the 44th Boston University Conference on Language Development,
ed. Megan M. Brown and Alexandra Kohut, 634-647. Somerville, MA: Cascadilla Press.

and the set denoted by the predicate (the set of dancers). (1) is true if the set of bunnies is a subset of the set of dancers. *Only* signals the opposite: (2) is true when the set of dancers is a subset of the set of bunnies. Given these set semantic similarities, will children who make errors with universal quantifiers make similar errors with adverbial quantifiers like *only*?

One of the most studied errors with universal quantification that children make are the so-called overexhaustive or spreading errors (see Philip 1995, Crain et al. 1996, Drozd 2001 and Geurts 2003 for overviews). Children incorrectly reject pictures like Figure 1 as a description of (1). These errors are frequent among children between the ages 4-9, peaking at 7 (Roeper et al. 2006, Aravind et al. 2017, de Koster et al. 2018), and are cross-linguistically robust (e.g. Drozd et al. 2019). When asked to explain what is wrong with the picture, they often point to other individuals who are also participating in the action as the reason for rejection, e.g. for (1) children would likely identify the dancing panda in Figure 1 as the reason for rejection.

Figure 1: Example of "All+1 Participant" picture type. All animals are depicted in the picture as participating in the action predicated by the sentence.

Figure 2: Example of "Two Participant" picture type. One of the mentioned animals and one additional animal are not participating in the action.

Children's explanations for rejecting such pictures suggests that they prefer a 1-to-1 pairing of individuals in the two sets. Such an interpretation could be described linguistically using the bi-conditional *all-and-only,* and set-theoretically children are requiring both sets to be the same set, i.e. all bunnies should dance, and all dancers should be bunnies. Some researchers have argued that these errors may in part be caused by a difficulty in determining which NP the quantifier is syntactically associated with in the sentence. Identifying this is crucial for figuring out the intended subset relation. For universal quantifiers, the NP specified by the quantifier needs to be a subset of the set denoted by the predicate.

When the adverbial quantifier *only* appears in pre-subject position in examples such as (2), the set denoted by the predicate needs to be the subset of the set denoted by the subject for the sentence to be true. If children find identifying this subset relation difficult, they might also reject pictures where some members of the set denoted by the subject do not participate in the action. For example, these children might reject pictures like Figure 2 with (2) because not all the bunnies are dancing, in essence treating *only* as if it meant *all-and-only.*

While there is previous work on children's knowledge of *only*, very little experimental research has investigated its properties as a quantifier relating two sets. *Only* is an adverb, so it is not part of a DP, and it can actually appear in any adverbial position.[1] Crain et al. (1994) looked at *only* in sentence initial position with single objects and subjects with sentences like *Only the kangaroo is holding a bucket*. They found that children aged 3-5 incorrectly accepted pictures where other animals in addition to the kangaroo were holding a bucket. These errors suggest that children may have interpreted these sentences as if *only* scoped over the object, e.g. *The kangaroo is holding only a bucket*, suggesting that 3-5 year old children have problems with correctly restricting the scope of pre-subject *only*.

But research with *only* modifying plural subject noun phrases has found that four year olds can correctly interpret *only*. Barner et al. (2011) found that 4 year old children correctly rejected intransitive sentences like *Only the dog and cow are sleeping* if, e.g. a cat was also depicted as sleeping.

Our investigations are instead focused on children's knowledge of *only* as a quantifier, examining whether or not children know that *only* does not require a 1-to-1 relationships between the sets it relates, e.g. *Only bunnies are dancing* is true even if not all the bunnies are dancing. As far as we know, this has not yet been studied. Because this interpretation relies on correctly identifying the subset relation signaled, we predict that there will be a correlation between making errors with *only* and making overexhaustive errors with universal quantifiers, because both those errors can be characterized as a failure to correctly identify a subset relation. We investigated this in three experiments.

3. Experiment 1: Picture Verification with *All* and *Only*

45 Dutch children (M_{age}=5;4, Range 4;4-6;4) took part in a 2 x 2 Picture Verification Task, with the factors PICTURE ("All+1 Participants", Figure 1, and "Two Participants", Figure 2) and QUANTIFIER (*All* vs. *Only,* using sentences like (3) and (4) in Dutch). 20 Dutch adults also participated (M_{age}=25).

(3) **Al** de konijnen dansen.
All the bunnies dance/are dancing.
(4) **Alleen** de konijnen dansen.
Only the bunnies dance/are dancing.

In the picture type "All+1 Participant" all members of the set mentioned performed the activity (e.g. dancing in Figure 1), with another animal character also participating in the activity (e.g. the dancing panda in Figure 1). In "Two Participant" pictures, one member of the named set did not participate in the action, and neither did an additional animal character (e.g. the panda in Figure 2).

[1] *Only* also associates with focus, but this aspect of its meaning is not relevant to the examples studied here, e.g. *Only BUNNIES are dancing* and *Only bunnies are DANCING*, do not differ in their requirements that no non-bunnies should dance, and more importantly, not all relevant bunnies need to participate for either sentence to be true.

Children saw 16 target items in total, four of each condition, plus three practice items and eight control items to check attention.

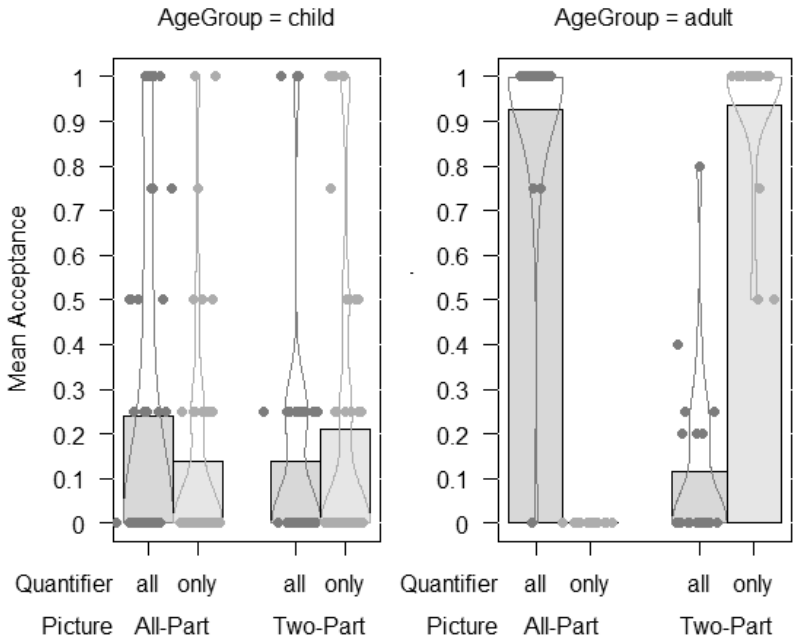

Figure 3: **RDI plot showing the results for Experiment 1.** Mean acceptances for all four conditions for children and adults are given as bars. Raw data is shown with dots representing participants, with horizontal jitter. Bean lines show the density of responses. While adults either accepted or rejected different conditions, with little variation, children showed a greater range of responses. Dots and the hour glass shape of the density lines show that for errors with the *all*-"All+1 Participant" condition a majority of the children either incorrectly rejected all items (overexhaustive/spreading errors) or accepted all items (adult-like). For the crucial *only*-"Two Participant" condition we see a similar distribution, with many children rejecting it, and many children showing adult-like acceptance.

3.1. Results Experiment 1

Children were nearly adult-like in their rejection of *only* sentences with "All+1 Participant" pictures and in their rejection of *all* sentences with "Two Participant" pictures. However, children incorrectly rejected the *all*-"All+1 Participant" pictures with *all*-sentences (76%), mentioning the dancing panda (overexhaustive or spreading error), and incorrectly rejected the "Two Participant" pictures with *only*-sentences (74%), mentioning the non-dancing bunny. These errors are consistent with interpreting both quantifiers as *all-and-only*.

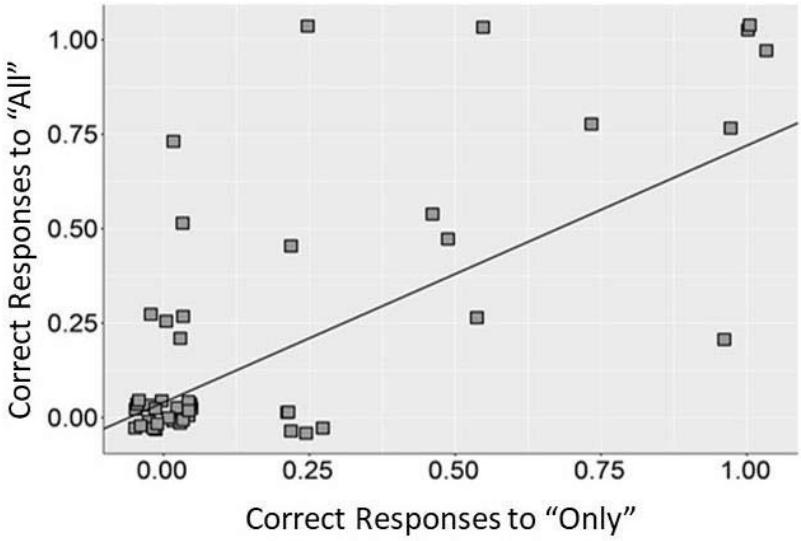

Figure 4: Results Experiment 1 Correlation between mean correct responses to *Only* (x-axis, Mean Correct answers to "Two Participant" items with *Only*) and mean correct responses to "All+1 Participant" pictures with *All* (y-axis) (Pearson coefficient: 0.72, $p < 0.001$, n=44)

We also analyzed children's responses using mixed effect models (Baayen et al. 2008). A model was preferred if its AIC value was at least two points lower. The maximal model that best explained the data retained the fixed factors *Quantifier* and *Picture* and the random factors *Participant* and *Item* only, with a significant interaction and without random slopes. Pairwise comparisons (Tukey) showed two significant interactions at the $p < 0.001$ level: the *all* ('*al*') - "All+1 Participant" picture condition (24% 'yes' responses) had significantly more 'yes' responses compared to the *only* ('*alleen*') - "All+1 Participant" picture condition (14% 'yes' responses) and the *all* ('*al*') -"Two Participant" picture condition (14% 'yes' responses) also had significantly fewer 'yes' responses compared to the *only* ('*alleen*') - "Two Participant" picture condition (26% yes responses). No other comparisons were significant.

We then compared children's errors incorrectly rejecting *only* with "Two Participant" pictures to their overexhaustive errors with *all* to see if there was a relationship between these errors. We found a significant positive correlation ($p<0.001$, Figure 4), meaning that errors with either quantifier correlate with errors with the other.

3.2. Discussion Experiment 1

We seem to have found clear indications that children's errors with universal quantifiers and *only* might be related. Children seem to incorrectly treat both quantifiers as if they meant *all-and-only.*

However, could children's incorrect interpretation of *only* be explained by other factors? One potential problem is the semantic contribution of the definite article used in Experiment 1, e.g. *All the bunnies* or *Only the bunnies,* see (3) and (4). The definite article implicates maximality in many cases. For example, in (5) we would expect all bunnies to dance, exhausting the participants. However, when the context indicates contrast, maximality can be suspended and the definite article seems to just pick out the relevant set, so (6) can be used to describe a situation where most of the girls are playing basketball, but the boys are all playing soccer.

(5) The bunnies are dancing.
(6) The girls are playing basketball.

Our adult participants did not interpret sentences with *only the X* as implicating that all the members of the set participated. However, our child participants might prefer to interpret the definite article maximally, which would make *only* equivalent with *all-and-only,* explaining their errors with *only*. In Experiment 2 we check to see whether children still make errors with *only* when no definite article is present, using sentences like (7) without a definite article.

(7) Alleen konijnen dansen.
Only bunnies dance/are dancing.

There is one additional issue. Without the definite article, sentences like (7) out of context suggest a generic statement, about e.g. all bunnies in the world. This is also because Dutch doesn't make a distinction between progressive and simple present/generic, using the same simple form for both. Even though the picture verification tasks itself strongly suggests that we are verifying the sentences solely based on the animals in the presented picture, it could still be an issue. To make sure that generic interpretations were also ruled out in Experiment 2, we created a frame story about a planet where there were several small groups of aliens. The aliens shown are the only ones of their kind, so statements about the aliens without a definite article could not be interpreted as generic.

4. Experiment 2: Picture Verification with *all* and *only* without definite articles

16 Dutch children (M_{age}=5;5, Range 4;5-6;11) took part in a 2 x 2 picture verification task, with the factors PICTURE ("All+1 Participants", Figure 5, and "Two Participants", Figure 6) and QUANTIFIER (*All* vs. *Only*, using sentences

like (8) and (9) in Dutch). 24 Dutch adults also participated as controls (M_{age}=24.38; Range 17-53).

(8) **Alle** kroepies dansen.
All kroepies are dancing.
(9) **Alleen** kroepies dansen.
Only kroepies are dancing.

All sentences had a pre-subject quantifier without a definite article, e.g. (8) and (9). The activities used were the same as in Experiment 1. We only changed the animals into novel fantasy creatures with non-existing names, e.g. Kroepies, Noppies, Blupsies, etc. and in the pictures we used novel creatures to represent the aliens (see Figures 5 and 6). To ensure that participants correctly interpreted the novel names, there was a training item before each target item where the names of both aliens in the upcoming target item were practiced.

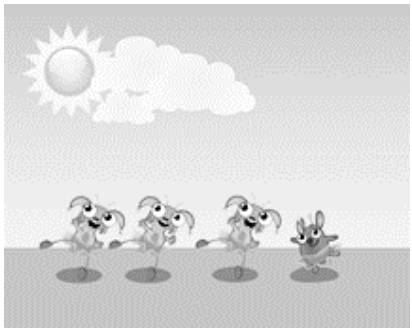

Figure 5. "All+1 Participant" picture type for Experiment 2. All aliens are participating in the action (dancing, show by the movements and tutu's).

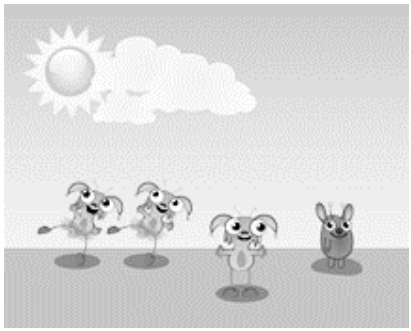

Figure 6. "Two Participant" picture type for Experiment 2. One of the mentioned aliens and one additional alien are not participating in the action (dancing).

4.1. Results Experiment 2

Figure 7 presents the percentage acceptance for all conditions for both the adults and the children. Adults show responses similar to Experiment 1. Children also show responses similar to Experiment 1 for *only* items. Children correctly rejected the *only*-sentences with the "All+1 Participant" pictures. But they still incorrectly rejected the *only*-"Two Participant" condition, indicating that removing the definite article did not lead to more adult-like interpretations for this condition.

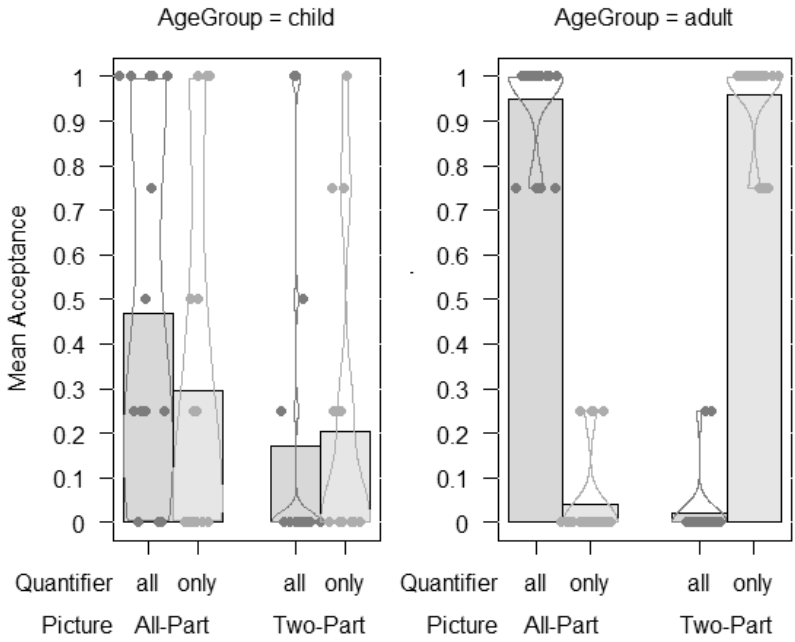

Figure 7. RDI plot of the results of Experiment 2 Children and adults, for *all* ('alle') and *only* ('alleen') with "All+1 Participant" pictures and "Two Participant" pictures. Bars show mean acceptance. Dots represent raw data for individual participants. Bean lines show the density of responses.

We analyzed responses using using mixed-effect models (Baayen et al. 2008). The AIC values were compared to determine which model best fit the data. The fixed effects were *Quantifier* (all vs. only), *Picture Type* ("All Participant" vs. "Two Participants"), and the random effects of Participant, Item and Age (in days, normalized), including random slopes (Jaeger 2008). The maximal model that explained significantly more variance retained the fixed factors Quantifier and Picture Type *only*, with a significant interaction. Pairwise comparisons were done (Tukey) on the children's results. Sentences with *All* with the "All+1 Participant" pictures were significantly less target-like than *All* with the "Two Participant" pictures ($p < 0.01$). Responses to *only* sentences were significantly more target-like with the "All+1 Participant" pictures than with the "Two Participant" pictures ($p < 0.001$). The condition *All*-"All+1 Participant" had a significantly higher accuracy than the condition *Only*-'Two Participants'. There was a significant difference between *All*-"All+1 Participant" and *only*-"Two Participant" ($p < 0.05$). No other differences were significant.

4.2. Discussion Experiment 2

The results of Experiment 2 show that children make the same errors regardless of the presence of the definite article. Removing the definite article did not lead to more adult-like responses for *only* sentences in the "Two Participant" condition. This strongly suggests that children's errors in rejecting *only* sentences in these cases is due to a non-adult like interpretation of *only,* and not due to a maximality requirement introduced by the definite article.

There is one additional issue. In both experiments we tested children with *all* and *only* at the same time. *All* modified the subjects, and required that the set denoted by the subject exhaustively participated in each action. We also used *only* in pre-subject position. This may have influenced children to prefer that the set denoted by the subject is always exhausted in the *only* items as well. One way to guard against an unwanted influence from universal quantifier items when testing *only* is to test each quantifier separately. We did just that in Experiment 3. Children's understanding of *only* was tested with a Situation Verification Task with Correction, and children's understanding of universal quantifiers were tested with a Picture Verification Task.

5. Experiment 3 Situation Verification with Correction

42 Dutch children (M_{age}=6.71, range 6;0-7;0) participated in two tasks, a Situation Verification with Correction Task and Picture Verification Task.

5.1. Situation Verification with Correction

In this task, multiple objects were distributed over six paper dolls, three girls and three boys, in four conditions (see Figure 8). Children were then asked if sentences with *only* like (8) correctly described the situation. This part of the task then gives us information similar to picture verification. If children rejected the sentences with the situation, they were then encouraged to fix the situation by moving or adding objects. This allowed us to systematically investigate why children rejected some items, and also to see what situation they would consider correct.

(8) **Alleen** meisjes hebben pizza.
Only girls have pizza.

	
Condition A True.	**Condition B** True.
	
Condition C False. Pizza slices in front of boys should be removed.	**Condition D** False. Pizza slices in front of boys should be removed.

Figure 8: **The four conditions presented in the Situation Verification Task in Experiment 3.** Three girl paper dolls are on the left, and three boy paper dolls are on the right, with pizza slices. In the table for each condition whether the situation is true or false is noted. Additionally, for the false situations, what needs to be done to correct the situation so that it would be true is also given.

We used four different situations (A-D), which corresponded to the four conditions (See Figure 8). Condition A was true, with all mentioned characters having the mentioned object. Condition B was also true, but in this case only one or two of the mentioned characters was given the object. This is the crucial condition to see if children reject *only* sentences because they want all members of the subject set to also be members of the set denoted by the predicate (having pizza). Conditions C and D were both false, because one of the non-mentioned characters also had the object. These conditions test if children know the basic meaning of *only,* i.e. that it excludes other characters from having the object. We expect children to reject these situations and correct them by removing objects from the non-mentioned character, e.g. the boys.

After participating in the Situation Verification task, all participants took part in a short Picture Verification Task.

5.2. Picture Verification with *each*

The picture verification task was used to identify overexhaustive (spreading) errors with universal quantifiers. We tested transitive sentences like (9) with the Dutch distributive quantifier *elke* (Eng. *each*) modifying the subject, with three different pictures: extra object pictures, extra subject pictures and correct 1-to-1 pictures, including practices and control items. Figure 9 presents an example of the materials.

(9) **Elke** jongen rijdt op een skateboard.
Each boy is riding a skateboard.

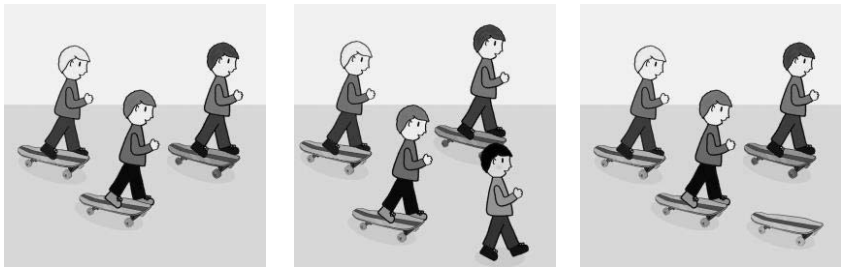

1-to-1 pairing **Extra subject** **Extra object**

Figure 9. Examples of the three types of pictures used in the Picture Verification Task in Experiment 3.

5.3. Results Experiment 3

For the Situation Verification Task, almost all children correctly accepted Condition A situations. Almost all children also correctly rejected Condition C and D situations. Children also correctly fixed the situations in C and D to make them true by removing objects from the non-mentioned character, e.g. removing pizza from the boys. These responses show that children know the basic meaning of *only*.

However, children incorrectly rejected and unnecessarily corrected Condition B situations 64% of the time, by e.g. giving all the girls pizza. This seems to be because they want all girls to exhaustively have pizza, and this response is what would be expected if children incorrectly treat *only* as *all-and-only*.

Similarly, even though children correctly rejected Condition D 87% of the time, because one of the non-mentioned characters had the object, in 69% of those rejections, they also unnecessarily gave objects to the mentioned character. For example, they would remove the pizza from the boys and give it to the girl who

did not have pizza, so that all girls had pizza. This correction also suggestions that children interpret *only* as *all-and-only*. Comparing the results on the Situation Verification Task with the Picture Verification Task that was used to test children's understanding of *all*, we again found a significant positive correlation between adult-like responses with *all* in the Picture Verification Task and with *only* in the Situation Verification (See Figure 10, Pearson coefficient, 0.33, $p < 0.05$, n=38).

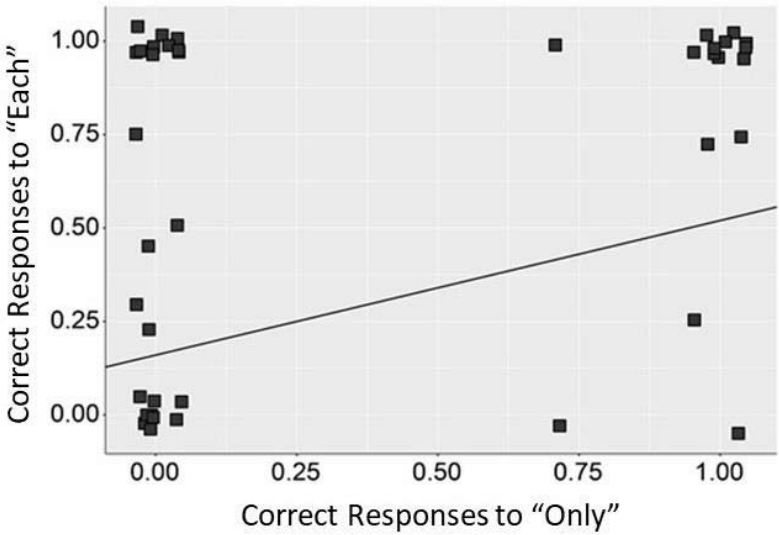

Figure 10: Results Correlation Experiment 3 Correlation between correct response to *Only* questions in Condition B in the Situation Verification Task (x-axis) and extra-object questions with *Each* in the Picture Verification Task (y-axis).

5.4. Discussion Experiment 3

In Experiment 3 we tested *Only* separately so that there would be no influence from universally quantified sentences. However, we still found the same errors we found in Experiments 1 and 2. Further, we again found a strong positive correlation between errors with *Only* and overexhaustive errors with *Each* in the Picture Verification Task.

Our results show that most children fell into three main categories which can be easily distinguished in Figure 10: (1) Adult-like (top-right corner) (2) *All-and-only* responses for both quantifiers (bottom-left) and (3) Adult-like with *Each* but give *all-and-only* responses for *only* (top left). Almost no children (n=3) gave correct answers for *only* but still made overexhaustive errors with *each*. This pattern of responses suggests that children first learn to correctly interpret universal quantifiers and then later learn to correctly interpret *only*.

6. General Discussion

In three experiments we investigated whether or not children who make overexhaustive errors with universal quantifiers also make similar errors with the adverbial quantifier *only*. In all three experiments we found that this was in fact the case, suggesting that these two errors are related. Both errors could be explained by children incorrectly treating universal quantifiers and *only* as if they meant *all-and-only*, rather than simply identifying the correct subset relation.

These two errors seem to be related. But can they be attributed to a similar cause? Consider briefly one of the main explanations for children's overexhaustive errors, the event quantification account (Philip, 1995, Roeper et al. 2006). This account claims that children's overexhaustive errors have to do with difficulties with identifying the scope of quantifiers. The youngest children treat universal quantifiers like sentential adverbs, quantifying over the entire sentence. Examine example (1) again, repeated here at (10):

(10) All the bunnies are dancing.

Under the event quantification account, sentences like (10) are interpreted as *Every event in the picture is a bunny dancing event.* Later, they learn that quantifiers syntactically associate with an NP, and that this association identify for e.g. universal quantifiers, which set is the subset (restrictor set). When they interpret the quantifiers as specifying the correct NP, they can then get an adult-like interpretation.

This account could be extended to also explain children's errors with *only*. If children first interpret sentences like (2) as *The only events are events of bunnies dancing*, they would reject cases where other animals dance and when some bunnies are not dancing. Like universal quantifiers, this account would predict children will stop making errors when they learn to restrict *only* correctly. However, because *only* does not syntactically combine with the NP it modifies, syntactic cues are not available for *only*. This analysis would then predict that *only* will take longer to learn. This prediction is consistent with what we found when looking at the correlations between results with universal quantifiers and *only.*

7. Conclusion

We conclude that (i) *only* errors and overexhaustive spreading errors are related (ii) children give responses consistent with interpreting both quantifiers as *all-and-only* to both quantifiers, giving responses suggesting the sets the quantifiers relate are the same, (iii) it seems that children either do not understand that these quantifiers signal subset relations, or find it difficult to identify the subset relations, and (iv) overexhaustive errors with universal quantifiers are overcome first, perhaps because syntactic structure more clearly identifies the subset (denotation of the N), because the quantifier syntactically specifies the NP.

References

Aravind, Athulya, Jill de Villiers, Peter de Villiers, Christopher Lonigan, Beth Phillips, Jeanine Clancy, Susan Landry, Paul Swank, Michael Assel, Heather Taylor, Nancy Eisenberg, Tracy Spinrad & Carlos Valiente (2017). Children's quantification with "every" over time. *Glossa: a Journal of General Linguistics* 2(1).

Baayen, Harald, Douglas Davidson & Douglas Bates (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59(4). 390–412.

Barner, David, Neon Brooks & Alan Bale (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition* 118(1). 84–93.

Barwise, Jon & Robin Cooper. (1981). Generalized quantifiers and natural language. In Philosophy, language, and artificial intelligence, 241–301. Springer.

Crain, Stephen, Rosalind Thornton, Carole Boster, Laura Conway, Diane Lillo-Martin & Elaine Woodams (1996). Quantification without qualification. *Language Acquisition* 5(2). 83–153.

Crain, Stephan, Ni, WJ and Conway, Laura (1994). Learning, parsing and modularity. In *Perspectives on Sentence Processing*, ed. C Clifton Jr., L Frazier, K Rayner, pp. 335–56. Dordrecht, Netherlands: Reidel

de Koster, Anna. M. B., Spenader, Jennifer., and Hendriks, Petra. (2018). Are Children's Overly Distributive Interpretations and Spreading Errors Related? Proceedings of the 42st Annual Boston University Conference on Language Development, 1:413-426.

Drozd, Kenneth (2001). Children's weak interpretations of universally quantified questions chap. in "Language Acquisition and Conceptual Development", M. Bowerman and S. Levinson, editors, 340–376. Cambridge University Press.

Drozd, K., Anđelković, D., Savić, M., Toskovic, O., Gavarró, A., Lite, A., Hrzica, G., Kovačević, M., Kraljević, J. K., Skordi, A., de López, K. J., Sundahl Olsen, L., Hollebrandse, B., van Hout, A., van Koert, M., Fabre, E., Hubert, A., Noveck, I., Ott, S., Yatsushiro, K., Balčiūnienė, I. & Ruzaite, J., A Crosslinguistic Study of Symmetrical Judgments, Proceedings of the 43rd annual Boston University Conference on Language Development.Brown, M. M. & Dailey, B. (eds.). Boston, USA: Cascadilla Press, Vol. 2. p. 217-230 14 p.

Geurts, Bart (2003). Quantifying kids. *Language acquisition,* 11(4). 197–218.

Philip, William Churchill. (1995). Event quantification in the acquisition of universal quantification: University of Massachusetts, Amherst, GLSA Publications dissertation.

Roeper, Thomas, Uri Strauss & Barbara Pearson (2004). The acquisition path of quantifiers: Two kinds of spreading. Current Issues in Language Acquisition, UMOP 34.

Proceedings of the 44th annual Boston University Conference on Language Development

edited by Megan M. Brown
and Alexandra Kohut

Cascadilla Press Somerville, MA 2020

ISSN 1080-692X
ISBN 978-1-57473-057-9 (2 volume set, paperback)

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