interpretation of plural-denoting expressions appears to diverge from the adult one until surprisingly late in their language development (Brooks & Braine, 1996; de Koster, 2021; Padilla-Reyes, 2018; Pagliarini et al., 2012).

In this paper, we investigate the potential connection between the distributivity/collectivity asymmetry attested in the adult language system and children's non-adult-like behavior. Specifically, with our study we aim to assess the hypothesis that in the adult language system the distributive reading of sentences with plural definite DPs is usually blocked via the generation of an implicature, and that children's well-attested difficulties with implicature generation can explain their non-adult-like interpretation pattern.

2. Background

2.1. Adults' interpretation of plural definites

As mentioned above, adults find the collective reading of sentences with plural definite DPs perfectly adequate. On the other hand, they appear to disprefer the distributive reading. When asked to judge sentences like (1) in combination with a distributive picture (Figure 2), adult participants' acceptance rate is invariably low: around 10% in Grinstead et al. (2021), 25% in the study of de Koster et al. (2018), and around 50% in Pagliarini et al. (2012). As argued in de Koster (2021), the observed variation in acceptance rate is likely caused by differences in the experimental methods.

How does this asymmetry between full acceptance of collective readings and a high rejection rate of distributive readings emerge in the adult language system? An intriguing explanation has been proposed by Dotlačil and colleagues (Dotlačil, 2010; Pagliarini et al., 2012; see also de Koster, 2021; Padilla-Reyes, 2018). We will refer to this explanation as the Implicature Account of distributivity/collectivity.

2.2. The Implicature Account of distributivity/collectivity

According to the Implicature Account, the interpretation of sentences with plural-denoting expressions (e.g., *the girls*, *all girls*) involves a competition between alternative, unpronounced forms, and unfolds in an analogous way to the generation of implicatures (e.g., the *some but not all*-implicature, Grice, 1975, 1989). Upon hearing a sentence such as (1), hearers have (at least) two interpretations available: the collective reading, and the distributive reading. The distributive reading, however, can unambiguously be expressed with a sentence like (2):

(2) Each girl is washing a dog.

Given that by uttering (1) the speaker chooses not to use a distributive marker (e.g., *each*), the hearer can conclude that the speaker's intended meaning corresponds to the collective reading (Figure 1), and not to the distributive reading

(Figure 2), because the latter can more straightforwardly and unambiguously be expressed with sentence (2).

The Implicature Account appears to fit the experimental data concerning adults' interpretation patterns. The distributive reading of sentences like (1) is not unacceptable, but strongly dispreferred because of a distributive/collective inference. What about the interpretation of plural-denoting expressions in children?

2.3. Children's interpretation of plural definites

Several studies show that children's interpretation of plural-denoting expressions appears to develop quite slowly (Brooks & Braine, 1996; de Koster, 2021; Padilla-Reyes, 2018; Pagliarini et al., 2012). This also holds for their interpretation of sentences with plural definite DPs. Consider for instance the aforementioned study of de Koster et al. (2018). In this study, adults and children were administered a truth-value judgment task. When presented with sentences like (1) in combination with a distributive picture (Figure 2), adults' rejection rate was 75%. On the other hand, the rejection rate of children between the ages of 4 and 8 was under 4%. At the ages of 9, 10, and 11 the rejection rates were 12%, 34%, and 40%, respectively (see Grinstead et al., 2021; de Koster et al., 2017, for comparable results). So, children show a non-adult-like pattern of interpretations until quite late in their language development.

How can children's non-adult-like patterns of interpretation be accounted for? If we assume, following the Implicature Account, that adults prefer the collective reading for sentences with plural definite DPs because they generate an implicature, then children's low rejection rate of sentences like (1) in combination with a distributive picture can straightforwardly be explained: this non-adult-like pattern can be reduced to difficulties with the process of implicature generation.

Children's difficulties with implicature generation, in fact, represent a well-attested finding in the pragmatic literature (e.g., Foppolo et al., 2012, 2020; Noveck, 2001; Papafragou & Musolino, 2003), with the majority of studies focusing on the *some but not all*-implicature. This implicature has been argued to emerge as follows. Upon hearing a sentence such as (3), mature hearers infer that the speaker's intended meaning corresponds to (4), even though the literal meaning of *some* is "at least one, and possibly all" and hence (3) is literally true also in a situation in which Charles watered *all* of his plants. The reason why (4) is the preferred interpretation for sentence (3) is linked to the existence of the non-pronounced but more informative alternative utterance (5). If hearers assume that speakers always try to provide as much information as possible (Grice, 1975, 1989), when the speaker utters sentence (3) and not sentence (5), the hearer can exclude the interpretation compatible with *all*. As a result, the meaning of *some* is pragmatically enriched with the negation of *all* and an implicature emerges. Thus, (4) is taken as the intended meaning of (3).

(3) Charles watered some of his plants.
(4) Charles watered some but not all of his plants.
(5) Charles watered all of his plants.

Importantly, children have been shown to struggle with this pragmatic inference: whereas adults tend to reject sentences like (3) in a context in which *all* plants have been watered, children tend to accept such sentences until the age of 5 or 6 (e.g., Foppolo et al., 2012). Therefore, there exists a clear parallel between the *some but not all*-implicature and the purported distributivity/collectivity implicature: both inferences seem to emerge in a similar way, both represent a challenge for children, and both take some time to develop.

In the next section, we describe an influential explanation of children's non-adult-like implicature generation rate, the Pragmatic Tolerance Account. We discuss its potential role in children's interpretation of plural-denoting expressions (Section 2.4). Subsequently, we describe the study we carried out to test the predictions emerging from the Implicature Account and the Pragmatic Tolerance Account (Section 3).

2.4. The Pragmatic Tolerance Account of implicature generation in children

According to the Pragmatic Tolerance Account (Davies & Katsos, 2010; Katsos & Bishop, 2011; Katsos & Smith, 2010), contrary to what is generally assumed children are able to distinguish between pragmatically felicitous and pragmatically infelicitous sentences with the quantifier *some*. This means that they can in fact recognize the suboptimal status of *some*-sentences in contexts in which *all* can be used. What makes children's performance non-adult-like is the fact that they are more tolerant than adults towards pragmatic violations. Sentence (3) in a context in which Charles watered all of his plants is surely underinformative and hence pragmatically infelicitous. However, it is not semantically false. According to the Pragmatic Tolerance Account, children's well-attested reduced rate of implicature generation emerges because of this greater tolerance towards pragmatic infelicity.

In line with this, Katsos and colleagues argue that the standard binary truth-value judgment task, routinely used in implicature studies, is not a sensitive measure of children's competence because it forces participants to either reject or accept sentences that are pragmatically infelicitous, but not false.

To provide experimental support to this hypothesis, Katsos and Bishop (2011, Experiment 2) carried out a ternary judgment task. Child and adult participants were presented with animations, and asked to judge the utterances of a fictional character, Mr. Caveman. In the critical condition, Mr. Caveman would utter a sentence such as "The mouse picked up *some* of the carrots", when in fact all of the carrots had been picked up by the mouse. Importantly, participants' task was not to reject or accept the utterance (as in the standard binary truth-value judgment task), but rather to reward Mr. Caveman using a 3-point scale. Specifically, they could choose a small, a big, or a huge strawberry, depending on their judgment of Mr. Caveman's utterance. Notably, whereas in the critical condition of binary

truth-value judgment tasks children usually accept, and adults usually reject, underinformative sentences, in Katsos and Bishop's (2011) ternary judgment task children expressed a strong preference (89%) for the middle option (a big strawberry), and so did adults. More in general, no differences emerged between the two groups. Katsos and Bishop (2011) interpret this result as evidence for the Pragmatic Tolerance Account: in binary judgment tasks, children's pragmatic tolerance masks their sensitivity to pragmatically infelicitous sentences. However, when tested with a more sensitive measure (e.g., a ternary judgment task) children perform adult-like.

Having presented the Implicature Account of adults' distributive and collective preferences and the Pragmatic Tolerance Account of children's implicature generation, let us now turn back to children's interpretation of plural-denoting expressions. If it is true that adults, when interpreting sentences with plural definite DPs, prefer the collective reading because they generate an implicature, then, in light of the Pragmatic Tolerance Account, we can expect children and adults to perform alike when tested using a ternary judgment task. In other words, when presented with sentences with plural definite DPs in combination with a picture representing the distributive reading (such as sentence (1) and Figure 2 above), we can expect both children and adults to show a preference for the middle option. This would be in line with the hypothesis that the adult interpretation of plural definite DPs involves a process of implicature generation (Implicature Account), and that children's non-adult-like behavior regarding implicatures is linked to their greater pragmatic tolerance (Pragmatic Tolerance Account).

We now describe our experiment, in which we aimed to test these hypotheses for the interpretation of plural definites using a ternary judgment task.

3. Methods

3.1. Participants

Fifty children (age range: 7:6 – 10;9, mean age: 8;10) and 20 adults (mean age: 22;1) took part in the experiment. All participants were native speakers of Dutch. Children were recruited via their primary school (located in the province of Overijssel, the Netherlands). The task was presented to both groups on the platform Qualtrics, but was carried out by children at school (individually, but under the supervision of a teacher), and by adults in a quiet location of their choice.

3.2. Materials and procedure

The materials were taken from de Koster et al. (2017). Our task included sentences with two different determiners (*de* 'the', *ieder* 'each') and pictures representing the two possible reading of sentences with these determiners (collective reading, distributive reading). Example sentences of the 4 conditions and the predicted responses for child and adult participants are shown in Table 1.

Participants saw 6 items per condition (24 items in total), plus 8 fillers (sentences that were unambiguously true or false and did not include plural expressions). The experiment also included 4 practice items.

At the beginning of the experiment, participants were introduced to a puppet, Benni the Icelandic dog. The participants' task was to help Benni learn Dutch. Benni would describe some pictures, and the participant could rate Benni's utterances using one star, two stars, or three stars. Participants were told that these options corresponded to 'wrong' (*fout*), 'a bit right' (*een beetje goed*), and 'right' (*goed*), respectively.

3.3. Predictions

For conditions *the*-collective and *each*-distributive, we expected both groups to show full acceptance and hence prefer the 3-star option. In fact, because of their lexical semantics, *the* is expected to be fully compatible with the collective reading and *each* with the distributive reading. As such, these two conditions can be considered control conditions. On the other hand, we expected our participants to reject *each*-collective items and hence to opt for the 1-star option in this condition: *each* is argued to be a distributive quantifier, so because of its lexical semantics a collective reading should not be available. Finally, in line with the Implicature Account and the Pragmatic Tolerance Account, we expected adults and children to choose the middle (2-star) option for *the*-distributive items.

Table 1. Experimental conditions and predictions

Condition	**Example sentence**	**Example picture**	**Predicted response (for adults and children)**
The-collective	De meisjes wassen een hond 'The girls are washing a dog'	Figure 1	3-star option
Each-distributive	Ieder meisje wast een hond 'Each girl is washing a dog'	Figure 2	3-star option
Each-collective	Ieder meisje wast een hond 'Each girl is washing a dog'	Figure 1	1-star option
The-distributive	De meisjes wassen een hond 'The girls are washing a dog'	Figure 2	2-star option

4. Results

Figure 3 shows adults' and children's percentages of response types in the four conditions.

In the two control conditions, *the*-collective (Figure 3, upper right panel) and *each*-distributive (Figure 3, lower left panel), as expected both adults and children showed a strong preference for the 3-star option. In the condition *the*-collective,

adults and children chose the 3-star option 99% and 90% of the time, respectively. Similarly, in the condition *each*-distributive, adults and children selected the 3-star option 99% and 95% of the time, respectively. Given these expected results, the control conditions will not be discussed further.

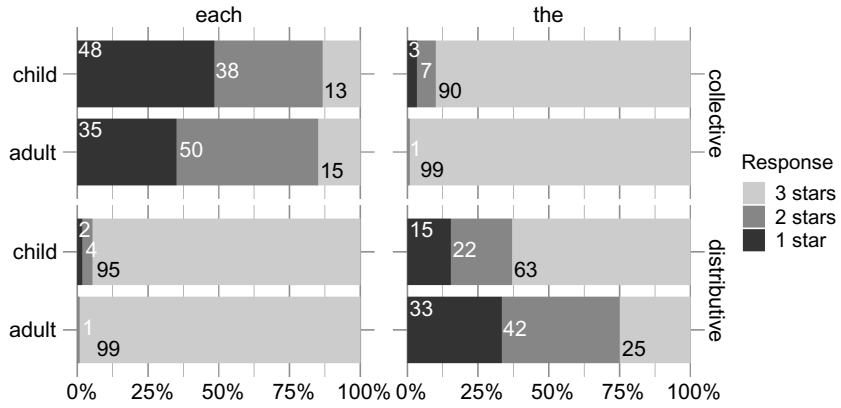

Figure 3. Percentages (rounded) of response type for the four conditions

The other two conditions, *the*-distributive and *each*-collective, were analyzed separately, fitting a series of ordinal regression models using the package *ordinal* (Christensen, 2019) in R (R version 4.0.4, RStudio Team, 2021). The proportional odds assumption was assessed using the graphical method recommended by Harrell (2001). Given that the predictor variable AGEGROUP violated this assumption, we refitted our models using the function *clmm2* of the same package, because *clmm2*, in contrast to *clmm*, allows for scale effects (although as a drawback it only allows for one random effect term). In a model with scale effects, the scale is allowed to differ between levels of the explanatory variable, therefore relaxing the proportional odds assumption (Christensen, 2018).

For both conditions, the best models included RESPONSE (an ordered factor with 3 levels: 1-star response, 2-star response, 3-star response) as outcome variable, AGEGROUP as predictor variable, and a random effect for PARTICIPANT. The goal of this analysis was to look for potential differences between the group of adults and the group of children.

In the *each*-collective model the variable AGEGROUP was not significant (β = -0.39, p = 0.49), indicating that children and adults behaved similarly.

On the other hand, in the *the*-distributive model AGEGROUP turned out to be a highly significant predictor (β = 4.01, $p < 0.01$), indicating a difference between children and adults. As we used the adult group as the reference level, the positive coefficient for AGEGROUP indicates that in the group of children a higher rating (2-star or 3-star response) was more likely than in the group of adults. To confirm that children's preference for a higher rating was driven by their preference for the 3-star rather than for the 2-star option (as Figure 3 already suggests), following the strategy described by Christensen (2015), we calculated the predicted

probabilities of the 2-star response and of the 3-star response for the two groups. According to our best model, children's and adults' predicted probability of selecting the 2-star option was 0.21 and 0.66, respectively. On the other hand, children's and adults' predicted probability of selecting the 3-star option was 0.76 and 0.08, respectively.

5. Discussion

According to the Pragmatic Tolerance Account put forward by Katsos and colleagues (e.g., Katsos & Bishop, 2011), children's apparent difficulties with implicature generation can be linked to their pragmatic tolerance. Despite showing non-adult-like performance on implicature generation in binary truth-value judgment tasks, children are predicted to show adult-like performance in more sensitive paradigms such as the ternary judgment task. Moreover, according to the Implicature Account of distributivity/collectivity (e.g., Pagliarini et al., 2012), the adult-like interpretation of plural definite DPs requires implicature generation. From these two accounts, it follows that ternary judgment tasks should allow children to perform adult-like when tested on their interpretation of plural definite DPs.

In this study, we experimentally assessed the validity of this prediction by testing children's and adults' interpretation using a ternary judgment task. Contrary to our expectations, despite the fact that this task is assumed to be more sensitive to pragmatic (in)felicity, children's interpretation of plural definite DPs was shown to differ from adults', exactly as in classic binary truth-value judgment tasks with plural definite DPs (e.g., de Koster et al., 2018). In our experiment, the option selected most often by children was the full acceptance option (the 3-star response), whereas the option selected most often by adults was the middle option (the 2-star response). Our results therefore seem at odds with the Pragmatic Tolerance Account and with the Implicature Account. Contrary to the predictions of these accounts, the children in our experiment did not perform adult-like, in spite of the pragmatically more sensitive experimental method used. Relatedly, the parallel between the *some but not all*-implicature and the distributivity/collectivity inference does not seem to hold, because children did not show a preference for the middle option (the 2-star response), as they do in ternary judgment tasks with classic scalar implicatures.

5.1. Restricting the Pragmatic Tolerance Account?

When trying to explain our unexpected results, a first issue that seems relevant is the generalizability of the Pragmatic Tolerance Account. When discussing their proposal, Katsos and Bishop (2011) appear to make predictions about quantity implicatures in general (i.e., scalar implicatures like the *some but not all*-implicature, and instances of *ad hoc* implicatures). However, despite suggesting that pragmatic tolerance may play a role also for other pragmatic phenomena, Katsos and colleagues (Davies & Katsos, 2010; Katsos & Bishop,

2011; Katsos & Smith, 2010) do not explicitly spell out the conditions under which children's pragmatically tolerant attitude should emerge. This raises the possibility that the Pragmatic Tolerance Account may hold for some but not all implicatures. In this regard, it is worth mentioning that recent studies on adults' implicature generation suggest that the behavior of implicatures (even of the same type) can be quite heterogeneous (Gotzner et al., 2018; Van Tiel et al., 2016). Therefore, it is conceivable that the quantity implicatures tested by Katsos and Bishop (2011) may not be representative of the whole category of implicatures, and that pragmatic tolerance may emerge only in connection with specific implicatures. As a matter of fact, the vast majority of acquisition studies of implicatures still focuses exclusively on the *some but not all*-implicature. In light of our results, it appears pivotal to address the issue of the generalizability of the Pragmatic Tolerance Account, especially using ternary judgment tasks (see Tieu et al., 2020 for an excellent example).

5.2. The alternative form in the Implicature Account

Before we turn to our unexpected results on the *the*-distributive condition in relation to the Implicature Account of distributivity/collectivity, let us first consider our findings regarding *each*. Recall that sentences with the distributive marker *each* are considered the optimal alternative for expressing distributive readings. Hence, their existence and interpretation is crucial for the purported implicature process through which sentences with plural definite DPs are interpreted non-distributively.

The findings of our experiment at first sight do not seem consistent with the assumption that Dutch *ieder* 'each' is a strong marker for distributivity. When presented with *each*-sentences in combination with a collective picture and given the possibility of choosing a response on a 3-point scale, our participants did not categorically opt for rejection: the 1-star response was only chosen 35% and 48% of the times by adults and children, respectively. Paradoxically, the children in our experiment seem to fit the predictions for *each* better than adults. However, as expected, neither children nor adults showed widespread acceptance of the collective reading of *each*, given that the 3-star response was chosen less than 15% of the times by our participants (see Figure 3).

One way to interpret these results is to concede that the quantifier we used (Dutch *ieder*) is more similar to English *every* than to English *each*. Despite being a distributive quantifier, the quantifier *every*, unlike *each*, can be used in situations of partial rather than full distributivity. Consider the following sentence (from Tunstall, 1998, p. 99):

(6) Jake photographed each/every student in the class.

Whereas *each* seems to be compatible only with a situation in which students were photographed individually (full distributivity), *every* appears felicitous also in a situation in which some of the photos portray more than one student (partial distributivity).

The hypothesis that Dutch *ieder*, exactly like English *every*, may be somewhat less marked for distributivity than English *each*, could explain participants' performance in the condition *each*-collective. This raises the question whether partial distributivity also has an influence on the generation of the distributivity/collectivity inference (and hence, on our results concerning the condition *the*-distributive). One could argue that if Dutch *ieder* is not always associated with the distributive reading, the distributivity/collectivity inference cannot consistently emerge. In other words, if sentence (7) is not interpreted exclusively with a distributive reading, it cannot represent a better alternative than (8) for expressing the distributive reading. Hence, the interpretation of (8) may not trigger, or may not require, implicature generation.

(7) Ieder meisje wast een hond.
'Each/every girl is washing a dog.
(8) De meisjes wassen een hond.
'The girls are washing a dog.

Despite these speculations, we do not believe that the possibility of partial distributivity for *ieder* represents a problem for the Implicature Account of distributivity/collectivity. Partial distributivity does not mean that the collective and the distributive reading are equally available. Rather, it means that, for the quantifier to be used appropriately, full distributivity is not always strictly required. Instead, at least in specific contexts (as illustrated by (6) above) some amount of distributivity suffices. This being said, it seems undeniable that in normal circumstances Dutch *ieder* is strongly associated with full distributivity, as is shown by the strong acceptance of *ieder* in distributive contexts in our experiment. Moreover, if we consider Champollion's (2019) distinction between simple universal quantifiers (e.g., *all*) and distributive universal quantifiers (e.g., *each*, *every*), it is clear that *ieder* must belong to the latter category. Like *each* and *every*, *ieder* is not compatible with predicates that refer to a collective action, such as *gather*:

(9) All philosophy professors gathered in the pub.
(10) * Each/every philosophy professor gathered in the pub.
(11) * Iedere filosofieprofessor verzamelde zich in de pub.

So *ieder* generally behaves like the English distributive quantifiers *each* and *every*. As such, in Dutch, sentence (7) is always a better alternative to convey the distributive reading compared to sentence (8).

In sum, on the one hand, we believe that the particular nature of Dutch *ieder* may be relevant for explaining our results in the condition *each*-collective, where, for both children and adults, we observed more 2-star responses than expected. On the other hand, we also believe that even if Dutch *ieder* was seen by our participants as partially distributive, this cannot have played a role in the condition *the*-distributive and does not yield an argument against the Implicature Account.

An alternative explanation for why the participants in our study did not show an absolute preference for the 1-star option in the condition *each*-collective may be related to the possibility that a collective action (as in Figure 1 above) can be conceived of as the sum of separate distributive actions (de Koster, 2021, Chapter 8). In this view, every girl is individually performing a "dog-washing" action ($girl_1$ is washing a dog ^ $girl_2$ is washing a dog ^ $girl_3$ is washing a dog) and the dog washed by each girl happens to be the same. If such an interpretation is indeed possible, but perhaps less preferred, language users may not categorically reject pictures representing a collective action like the one in Figure 1 for *each*-sentences. This could arguably explain the high rate of 2-star responses in condition *each*-collective of our experiment: participants recognized the distributive nature of the *each*-sentences, but perhaps judged them as not completely inappropriate in a situation in which there is a single dog that is being washed by different girls (Figure 1). Syntactically speaking, such an interpretation for sentence (7) can in fact be derived if language users allows for an inverse scope reading (with *a* taking scope over *each*, see Champollion, 2019). Importantly, however, an inverse scope reading appears to be strongly dispreferred for both *every* and *each* (Feiman & Snedeker, 2016), and this may be another reason why in our experiment we found a preference for 1-star and 2-star responses, rather than for 3-star responses. As an in-depth discussion of this issue goes beyond the scope of this paper, we leave this for further research.

5.3. Is the distributivity/collectivity inference really a scalar implicature?

Returning to the condition *the*-distributive, we now address a fundamental issue related to the Implicature Account, namely whether the distributivity/collectivity inference can in fact be considered a scalar implicature in the first place.

In their formulation of the Implicature Account, Pagliarini et al. (2012) explicitly characterize the distributivity/collectivity inference as a conversational implicature. The implicature generation process they illustrate, however, seems to describe specifically a scalar implicature. A scalar implicature view is indeed unambiguously endorsed by Padilla-Reyes and colleagues (Grinstead et al., 2021; Padilla-Reyes, 2018).

Scalar implicatures (with the classical *some but not all*-implicature as the prototypical example) are a type of quantity implicature; quantity implicatures, in turn, are a type of conversational implicature. What characterizes scalar implicatures is the fact that they are triggered by lexical scales (sometimes called *Horn scales* after Horn, 1972, 1989), that are defined by asymmetrical entailment. For instance, in the scale <*some*, *all*>, the stronger term *all* is more informative than the weaker term *some*, so the quantity of information expressed by the former is higher.

Does the distributivity/collectivity inference meet the criteria to be classified as a scalar implicature? In other words, in what sense are *each*-sentences more informative than *the*-sentences?

To answer this question, of particular relevance is the existence of a class of scalar implicature triggers called privative dyads (Horn & Abbott, 2012). Consider the privative dyad <*rectangle*, *square*>. In this dyad, the stronger term *square* is marked for a particular feature (equilaterality) for which the weaker term *rectangle* is not marked. Thus, in normal circumstances, the meaning of *rectangle* is enriched with the negation of the property in question, and we generally think of rectangles as being non-equilateral. Arguably, the same reasoning can be applied to the difference between the *each*-sentence in (7) and the *the*-sentence in (8): the former is marked for a particular feature (distributivity), for which the latter is not.

In essence, the similarities between the distributivity/collectivity inference and scalar implicatures, especially those based on private dyads, seem undeniable. Can we reconcile this observation with our data, which show that the distributivity/collectivity inference does not fully pattern with the *some but not all*-implicature? Bearing in mind the results from the conditions *each*-collective and *the*-distributive, we suggest a tentative explanation. This explanation involves another type of conversational implicature, that shares many features with scalar and quantity implicatures but is linked (exactly as the distributivity/collectivity implicature appears to be) to the concept of markedness: the manner implicature. Manner implicatures can arise when a marked (e.g., morphologically or syntactically more complex and less common) form becomes associated with a more marked (e.g., more specific and non-stereotypical) meaning. Now, a sentence with a distributive marker (e.g., *each*) differs in its form from a sentence with a plural definite DP (see (8) above). The former type of sentence can be considered more marked: it is less common, and it is syntactically more complex, especially when *each* is used adnominally or adverbially (*The girls are washing a dog each*). Therefore, despite the fact that perhaps *each*-sentences are not entirely unacceptable in collective contexts (as seen in Section 5.2, we can conceive a collective action also as the sum of different distributive actions), they generally become associated via a manner implicature with interpretations that are more marked because they express the purely distributivity reading (Figure 2). Sentences with *the*, on the other hand, become associated with unmarked collective interpretations.

Pragmatic tolerance has not been claimed to be a factor in the generation of manner implicatures. This, perhaps, could explain why the similarities between distributivity/collectivity inferences and scalar implicatures seem to break down in a ternary judgment task, and a more complex acceptance pattern emerges for distributivity/collectivity inferences.

The idea that the distributivity/collectivity inference may be an instance of manner implicatures remains speculative, but opens the door to new explorations of the interpretation of plurality.

References

Brooks, Patricia J., & Braine, Martin D. S. (1996). What do children know about the universal quantifiers all and each? *Cognition*, *60*(3), 235–268. https://doi.org/10.1016/0010-0277(96)00712-3

Champollion, Lucas. (2019). Distributivity in Formal Semantics. *Annual Review of Linguistics*, *5*(1), 289–308. https://doi.org/10.1146/annurev-linguistics-011718-012528

Christensen, Rune Haubo B. (2015). *A Tutorial on fitting Cumulative Link Mixed Models with clmm2 from the ordinal Package*. https://cran.r-project.org/web/packages/ordinal/vignettes/clmm2_tutorial.pdf

Christensen, Rune Haubo B. (2018). Cumulative link models for ordinal regression with the R package ordinal. *Submitted in J. Stat. Software*. https://mran.microsoft.com/snapshot/2020-02-24/web/packages/ordinal/vignettes/clm_article.pdf

Christensen, Rune Haubo B. (2019). *ordinal—Regression Models for Ordinal Data. R package version 2019.12-10.* https://cran.r-project.org/package=ordinal

Davies, Catherine, & Katsos, Napoleon. (2010). Over-informative children: Production/comprehension asymmetry or tolerance to pragmatic violations? *Lingua*, *120*(8), 1956–1972. https://doi.org/10.1016/j.lingua.2010.02.005

de Koster, Anna M. B. (2021). *Acting Individually or Together? An Investigation of Children's Development of Distributivity*. Ph.D. thesis, University of Groningen. https://doi.org/10.33612/diss.169356700

de Koster, Anna M. B., Dotlačil, Jakub, & Spenader, Jennifer. (2017). Children's understanding of distributivity and adjectives of comparison. *Proceedings of the 41st Annual Boston University Conference on Language Development*, 373–386.

de Koster, Anna M. B., Spenader, Jennifer, & Hendriks, Petra. (2018). Are Children's Overly Distributive Interpretations and Spreading Errors Related? *Proceedings of the 42nd annual Boston University Conference on Language Development*, 413–426.

Dotlačil, Jakub. (2010). *A study of same , different , reciprocals and others*. Ph.D. thesis, Utrecht University.

Feiman, Roman, & Snedeker, Jesse. (2016). The logic in language: How all quantifiers are alike, but each quantifier is different. *Cognitive Psychology*, *87*, 29–52. https://doi.org/10.1016/j.cogpsych.2016.04.002

Foppolo, Francesca, Guasti, Maria Teresa, & Chierchia, Gennaro. (2012). Scalar Implicatures in Child Language: Give Children a Chance. *Language Learning and Development*, *8*(4), 365–394. https://doi.org/10.1080/15475441.2011.626386

Foppolo, Francesca, Mazzaggio, Greta, Panzeri, Francesca, & Surian, Luca. (2020). Scalar and ad-hoc pragmatic inferences in children: guess which one is easier. *Journal of Child Language*, 1–23. https://doi.org/10.1017/S030500092000032X

Gotzner, Nicole, Solt, Stephanie, & Benz, Anton. (2018). Scalar Diversity, Negative Strengthening, and Adjectival Semantics. *Frontiers in Psychology*, 9(SEP), 1–13. https://doi.org/10.3389/fpsyg.2018.01659

Grice, Herbert Paul. (1975). Logic and Conversation. In *Speech Acts*. Brill. https://doi.org/10.1163/9789004368811_003

Grice, Herbert Paul. (1989). *Studies in the Way of Words*. Harvard university press.

Grinstead, John, Padilla-Reyes, Ramón, & Nieves-Rivera, Melissa. (2021). A Collective-Distributive Pragmatic Scale and the Developing Lexicon. *Language Learning and Development*, *17*(1), 48–66. https://doi.org/10.1080/15475441.2020.1863808

Harrell, Frank E. (2001). *Regression modeling strategies: with applications to linear models, logistic regression, and survival analysis*. Springer.

Horn, Laurence R. (1972). *On the semantic properties of the logical operators in English*. Ph.D. thesis, UCLA, Los Angeles.

Horn, Laurence R (1989). *A natural history of negation*. University of Chicago Press.

Horn, Laurence R., & Abbott, Barbara. (2012). < the, a>: (In)definiteness and Implicature. In *Reference and referring*. 325–356. MIT Press.

Katsos, Napoleon, & Bishop, Dorothy V. M. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition*, 67–81. https://doi.org/10.1016/j.cognition.2011.02.015

Katsos, Napoleon, and Smith, Nafsika. (2010). "Pragmatic Tolerance or a Speaker-Comprehender Asymmetry in the Acquisition of Informativeness?" in *Proceedings of the 34th Annual Boston University Conference on Language Development*, 221–232.

Noveck, Ira A. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition*, *78*(2), 165–188. https://doi.org/10.1016/S0010-0277(00)00114-1

Padilla-Reyes, Ramón. (2018). *Connections Among Scales, Plurality, and Iintensionality in Spanish*. Ph.D. thesis, The Ohio State University.

Pagliarini, Elena, Fiorin, Gaetano, & Dotlačil, Jakub. (2012). The acquisition of distributivity in pluralities. *Proceedings of the Annual Boston University Conference on Language Development*, *2*, 387–399.

Papafragou, Anna, & Musolino, Julien. (2003). Scalar implicatures: Experiments at the semantics-pragmatics interface. *Cognition*, *86*(3), 253–282. https://doi.org/10.1016/S0010-0277(02)00179-8

RStudio Team. (2021). *RStudio: Integrated Development Environment for R*. RStudio, PBC. http://www.rstudio.com/

Tieu, Lyn, Bill, Cory, Romoli, Jacopo, & Crain, Stephen. (2020). Testing theories of plural meanings. *Cognition*, *205*(August), 104307. https://doi.org/10.1016/j.cognition.2020.104307

Tunstall, Susanne Lynn. (1998). *The interpretation of quantifiers: Semantics and processing*. Ph.D. thesis, University of Massachusetts Amherst.

Van Tiel, Bob, Van Miltenburg, Emiel, Zevakhina, Natalia, & Geurts, Bart. (2016). Scalar Diversity. *Journal of Semantics*, *33*(1), 137–175. https://doi.org/10.1093/jos/ffu017

Proceedings of the 46th annual Boston University Conference on Language Development

edited by Ying Gong
and Felix Kpogo

Cascadilla Press Somerville, MA 2022

ISSN 1080-692X
ISBN 978-1-57473-077-7 (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