Quantifier spreading and the question under discussion

Dimitrios Skordos,^{a,b} Allyson Myers,^c David Barner^a

^a University of California, San Diego

^b University of Calgary

^c San Diego State University

ABSTRACT

Children often display non-adult-like behaviors when reasoning with quantifiers and logical connectives in natural language. A classic example of this is the symmetrical interpretation of universally quantified statements like “Every girl is riding an elephant”, which children often reject as false when they are used to describe a scene with, e.g., three girls each riding an elephant and a fourth elephant without a rider. We present evidence that children’s understanding of these sentences is not attributable to syntactic, semantic, or general processing limitations. Instead, in two experiments, we argue that children’s behavior stems primarily from difficulty in correctly identifying the speaker’s intended “question under discussion”, and that when this question is made contextually unambiguous, children’s judgments are almost completely adultlike.

Keywords: Inference, access to alternatives, class inclusion, quantifier spreading, pragmatic development, question under discussion, QUD

Access to alternatives and universal quantification in language acquisition

Introduction

Natural language affords humans the ability to represent and reason with logical connectives that encode conjunction (*and*), disjunction (*or*), the conditional (*if... then*), biconditional (*if and only if... then*), and negation (*it is not the case that...*), as well as existential and universal quantifiers (*some, every*). While these logical capacities are clearly required for language, there is significant controversy regarding their developmental origin and the relative contributions of innate primitive structures and learning. Although some recent studies argue that logical capacities are present in human infants even before they begin to speak (Cesana-Arlotti, Martín, Téglás, Vorobyova, Cetnarski, & Bonatti, 2018; Cesana-Arlotti, Kovács, & Téglás, 2020) work, beginning with Piaget (e.g., Inhelder & Piaget, 1964), has documented many different ways in which children's logical language differs from that of adults, sometimes indicating a protracted developmental sequence (e.g., Austin et al., 2014; Feiman, Mody, Sanborn, & Carey, 2017; Mody & Carey, 2016; Morris, 2008; Horowitz, Schneider, & Frank, 2018; Papafragou & Musolino, 2003; Singh, Wexler, Astle-Rahim, Kamawar, 2016).¹ Such findings, however, leave open the nature of children's difficulties, and whether they are attributable to immature logical representations, the problem of mapping existing logical representations to linguistic form, or instead to pragmatic or extra-linguistic factors like processing limitations that restrict the ability to express knowledge (Freeman et al., 1982; Crain et al., 1996; Geurts, 2003; Hollebrandse, 2004; Brooks & Sekerina, 2005; Drozd & van Loosbroek, 2006).

¹ Many studies document differences between adults and children, and we therefore select a small number of more recent findings from several distinct literatures.

Figure 1. A stimulus example like the ones used by Inhelder and Piaget (1974)

One classic example of such protracted logical development, first investigated by Inhelder and Piaget (1964), concerns children’s symmetrical interpretation of universally quantified expressions during the preschool and early elementary school years (see also Piaget, 1965; Inhelder & Piaget, 1974), sometimes called “quantifier spreading” (Roeper & de Villiers, 1993). In their work on this topic, Inhelder and Piaget showed a group of 4- to 7-year-old children a set of red squares as well as sets of red and green circles (see Figure 1), and then asked them, e.g., “Are all of the squares red?”, to which they often responded “no”, based on the fact that some of the circles were also red. In related work, Donaldson and Lloyd (1974) found that 4-year-old children treated the sentence, “All the cars are in the garages” as entailing that “All the garages have cars in them” indicating a symmetry requirement (see also Donaldson & McGarrigle, 1974, for a more widely available report of this phenomenon). Since these early studies, many additional studies have investigated this phenomenon, positing a range of semantic, pragmatic, and syntactic explanations, many as recent as the past 2-3 years (for very recent ideas, see Brunetti et al., 2021; Chen, Rosenstein & Hackl, 2020; Denić & Chemla, 2020; Drozd et al., 2019; Ke & Gao, 2020; Kiss, & Zétényi, 2017; Sekerina, & Sauermann, 2017; see Brooks & Parshina, 2019 for review). In the present study, we present evidence in favor of a

form of pragmatic account, and show that children's apparent errors arise when experimenters fail to adequately specify a clear "Question Under Discussion" (henceforth QUD; Roberts, 1996/2012). Thus, we argue that sensitivity to pragmatic violations, rather than a lack of logical representations, best explains children's symmetrical interpretation of universal quantifiers.

Background

Previous accounts of children's non-adult logical reasoning can be broadly classified into two types. The first class of accounts, expressed by Piaget and colleagues, posits that difficulties are either syntactic or semantic in nature. According to such accounts, children do not yet have the relevant formal representations necessary for adult-like computations, resulting in developmental discontinuity. These representations are gradually constructed over development, thus explaining children's early difficulty with cases like universal quantification. The second class of accounts generally assumes no qualitative differences in the syntactic or semantic representations of children and adults, and instead argues for continuity. On these accounts, although children have the relevant representations from early in acquisition, non-adult judgments arise due to either pragmatic limitations, or domain-general cognitive limitations – e.g., related to processing capacity or working memory. On this view, children may have the relevant logical "competence" but simply not have the ability to express their knowledge, due to "performance" limitations.

Versions of each hypothesis have been posited to explain children's symmetrical interpretation of universal quantifiers. Some researchers, beginning with Donaldson and Lloyd (1974), have argued that children lack the syntactic knowledge needed to evaluate complex quantified expressions, and that they therefore rely instead on contextual information like salience (see Smith, 1980, for discussion). For example, in a scene like the one in Figure 2, when

children hear “Every girl is riding an elephant” they may fail to appreciate the syntactic relationship between the quantifier “every” and the remainder of the sentence, and thus infer that the salient omission of a girl on the fourth element must be relevant, leading them to judge the sentence as false.

Figure 2. Girls riding elephants

Providing a more developed version of this hypothesis, others have proposed that children’s difficulties stem their failure to restrict the application of quantifiers in the grammar: Children’s non-adult-like judgements arise because they interpret the universal quantifier *every* as applying at both to the NP in which it is contained, but also to the NP that is the object of the verb (Roeper & de Villiers, 1993). For example, in the sentence, “Every girl is riding an elephant”, unlike adults, children do not restrict quantification as applying exclusively over girls, but rather “spread” its application toto elephants as well, effectively interpreting the sentence as “every girl is riding an elephant and every elephant is ridden by a girl”. Others have argued that errors could arise because children quantify over events as opposed to entities (Philip, 1991; 1995), or otherwise misidentify the syntactic features of the quantifier (Kang, 2001). Finally, a third group of researchers (Drozd & van Loosbroek, 1999; Drozd, 2001) have suggested that children might have a different semantics for the universal quantifier; errors may arise because

they treat *every* as a “weak” quantifier, similar to *many*, and therefore allow its interpretation to be determined contextually (Westerståhl, 1985). For example, consider the sentence in (1a) and its possible interpretations in (1b) and (1c):

- (1) a. Many Americans crew the International Space Station.
- b. A large proportion of the American population crews the ISS.
- c. A large proportion of the ISS crew are Americans.

One reading of (1a) corresponds to (1b), which, in most other cases would be the preferred reading of *many*, is known to be false and is therefore dispreferred (it is not actually the case that a large number of the set of US citizens are astronauts that crew the ISS). Instead, people judge (1a) to be true with a reading like (1c), effectively allowing the quantifier *many* to combine with the verb phrase of the sentence (“... crew the ISS”) as opposed to the subject. In this case, adult speakers appear to disregard the canonical compositional meaning of a quantified sentence and pragmatically infer an alternative speaker meaning (cf. ‘Reverse Interpretation’; Cohen, 2001; see Clothier, 2019 for recent experimental work on this question). On analogy, children, unlike adults, may treat “every” as a weak quantifier. When presented with a sentence like, “Every girl is riding an elephant”, they may allow “every” to quantify over “elephant”, resulting in a logical meaning akin to “Every elephant is ridden by a girl”.

In contrast to these theories that posit syntactic or semantic deficits, proponents of grammatical continuity have appealed to evidence that children are essentially adult-like in their representation of logical language (e.g., Gualmini, 2014; Gualmini, Meroni, & Crain, 2003), and have argued that their non-adult behaviors are due to either pragmatic differences or domain-general factors, such as limits to cognitive resources like working memory. Geurts (2003), for instance, argues that although children have an adult-like semantics of *every*, processing

difficulties associated with quantified expressions lead them to effectively treat *every* as semantically weak. Specifically, because processing limits make them unable to use the syntax of utterances to restrict the NPs to which quantifiers apply, children instead rely on pragmatic and contextual information, with the result that they allow *every* to quantify over *elephant* (or equivalent NPs) when they are focused contextually (e.g., as they are when an extra elephant is presented in a scene). Consistent with a role for domain general factors, even adults sometimes make child-like quantifier spreading errors when under cognitive load, despite having ostensibly mature syntactic and semantic representations (Bott & Shlotterbeck, 2018; Brooks & Sekerina, 2006; Meroni et al., 2006). Similar arguments have suggested that children may experience difficulty suppressing attention to the salient extra object in the scene (Freeman et al., 1982; cf. Brooks & Sekerina, 2005; Gouro, Norita, Nakajima & Ariji, 2001; Minai et al., 2012; O’Grady et al., 2010; Sekerina & Sauermann, 2015; Sekerina et al., 2018; Sugisaki & Isobe, 2001) and that more carefully designed experiments lead to more adult-like judgments. Finally, some have argued that children have difficulty restricting the domain of discourse to the relevant set referred to by the quantified subject of the sentence (Drozd & van Loosbroek, 2006; cf. Hollebrandse, 2004). For example, children may think that because one elephant lacks a rider, there must be a missing fourth girl who is not riding her elephant, making it true that not every girl is riding an elephant.

Germane to the present study, Crain and colleagues (Crain et al., 1996) proposed a form of pragmatic explanation for children’s judgments. In particular, they argued that children provide non-adult judgments because the question they are asked is pragmatically infelicitous, forcing them to entertain other possible questions that might make better sense given the context. Specifically, they argued that, “it is appropriate to ask a Yes/No question only if it is plausible to

entertain both ‘Yes’ and ‘No’ answers” (p. 114), an idea that they call *plausible assent/dissent*. Given this, a statement like, “Every girl riding an elephant” is infelicitous in the contexts used in quantifier spreading studies because the assertion is obviously true. The problem, according to Crain et al., is that in such cases, no alternative other than the actual outcome is ever under consideration, and therefore children have no basis for doubting the truth of the statement - a condition Crain et al. call “plausible dissent”. Consequently, children assume that the speaker must have some other claim in mind, and converge upon a statement that *does* satisfy the condition of plausible dissent, such as “Every elephant is ridden by a girl”.

Based on this logic, Crain et al. argued that when the conditions of plausible dissent are met, and the truth of the assertion is in doubt at some point during the experiment, children should view the assertion as felicitous and accept it as true. Take, for example, the situation in (2), adapted from Crain et al.:

(2) Background: Every girl is riding <some animal>.

Assertion: (Every girl is riding) an elephant.

Possible Outcome: Some girl is riding a dinosaur.

Actual Outcome: Every girl is riding an elephant.

In such a context, Crain et al. argue that children should readily accept the assertion that “Every girl is riding an elephant” since the possible outcome that some girl is riding a dinosaur licenses the generation of an alternative assertion which can be endorsed in its place – i.e., the assertion involving dinosaurs. For Crain et al., the availability of a plausible alternative provides a basis for felicitously evaluating the original statement.

As evidence for the hypothesis, Crain et al. presented children with scenarios in which plausible dissent was satisfied and found that under such conditions children accepted

universally quantified statements. For example, in one scenario presented to children, three characters considered drinking either sodas or cider, and ultimately decided to drink cider, making the statement “Every girl drank cider” true, but also felicitous (since the girls could have chosen sodas). Based on this success, Crain et al. argued that children therefore deployed an adult-like semantics, and did not make spreading errors. However, while the conditions of plausible dissent created in their paradigm rendered the utterances felicitous and allowed children to imagine a world in which the test statements might be false, it’s not clear that these judgments were based on adult-like grammatical representations of the sentences, or that the spreading phenomenon was actually eliminated. This is because in the Crain et al. study focus was placed on the choice between kinds of drinks, such that plausible dissent could be satisfied by attending purely to which drink was chosen and not to how many drinks were consumed. If children answered the experimenter’s question on the basis of kinds of things, rather than quantification, we would expect adult-like judgments even if children completely ignored quantification.

One framework for understanding this issue is Roberts’ notion of the “Question Under Discussion” or QUD (Roberts, 1996/2012). On this view, a QUD is defined as a set of alternative propositions that correspond to a topic set out in a discourse or conversation (see also Beaver, Roberts, Simons, & Tonhauser, 2017; Simons, Tonhauser, Beaver, & Roberts, 2010). For any particular surface string, such as ‘Every girl is riding an elephant’, multiple QUDs are possible, since different alternatives may be relevant to the topic at hand. For example, if the speaker is interested in what kinds of animals a girl might ride, then alternatives might include statements that differ with respect to kinds of animals, as in (3a). However, the speaker might also have a different QUD in mind, which generates alternatives that contrast with respect to the number of

girls riding elephants, as in (3b). Many other forms of alternative are also possible (e.g., 3c, 3d). Corresponding to this, not only do different questions entail different alternatives, but different alternatives entail different QUDs. For example, if alternatives that mention different kinds of animals are proffered to a listener, they should infer that the speaker is interested in what kinds of animals are being ridden, whereas if alternatives contrast quantificationally, they should infer that the speaker is interested in the number of animals being ridden.

- (3) a. Every girl is riding a *dinosaur*.
 b. *Some* of the girls are riding elephants.
 c. Every *boy* is riding an elephant.
 d. Every girl is *feeding* an elephant.

This notion of QUD is important to understanding quantifier spreading because it highlights the fact that a single response to a statement - e.g., accepting or rejecting a universally quantified statement - can in fact represent the endorsement of many distinct propositions. Consequently, although plausible dissent can be satisfied in ways that are germane to the speaker's intended QUD, a basis for rejecting a statement can also be provided that is not related to the intended question and that therefore doesn't address the spreading phenomenon. For example, in the Crain et al. study, the experimenter is interested in how many girls ordered cider, but the alternatives on offer are also compatible with a distinct QUD that relates to the kind of drinks the girls order - soda or cider. Thus, when children in the experiment see soda or cider considered, they could infer a corresponding QUD and base their response on the kind of drink. In such a scenario, their endorsement of the universally quantified expression might have little to do with an adult-like interpretation of the quantifier, and might instead be based wholly on their concern with sodas and ciders.

Lending credence to this concern, Skordos and Papafragou (2016) argued that children's well-attested failures to compute some/all scalar implicatures arise not because children can't compute implicatures, but because they evaluate quantified expressions on the basis of alternative QUDs that are relevant to the kinds of things described, rather than their quantity. To demonstrate this, Skordos and Papafragou showed that children could compute implicatures so long as other statements in the experimental context contrasted according to quantity, but not according to the kinds of things mentioned. In the study, children were tested with uninformative uses of the word *some* as in other studies of implicature by asking them to judge the felicity of a statement like "Some of the blickets have a scarf". However, before these test trials, children were first tested with a block of trials including the quantifier *all*, which differed crucially between two groups of subjects. Specifically, in the 'quantity' condition, children heard statements like "All of the blickets have a scarf" in a context where only 3 out of 4 blickets had a scarf and the fourth one had nothing, whereas in the 'object' condition children heard the same statement in a context where false because all four blickets had a different kind of object instead (e.g., a shovel). Skordos and Papafragou hypothesized that, when tested with a critical *some* statement like "Some of the blickets have scarves", children in the object group would base their judgements on the kind of object mentioned (e.g., contrasting whether the blickets have scarves vs. shovels), whereas children in the quantity group would base their judgments on the strength of the quantifiers (e.g., contrasting whether some or all of blickets have scarves). In keeping with their hypothesis, they found that preschoolers could easily derive scalar implicatures with the quantifier *some* in the quantity group, but did not do that in the object group.²

² It should be noted here that adults do not seem to be immune to such effects: In fact similar effects of the QUD on the generation of scalar implicatures have been documented in the literature (see Zondervan, 2007; Zondervan et al., 2008).

The findings of Skordos and Papafragou are important for two reasons. First, they show that when children interpret utterances according to a non-quantificational QUD, they may fail to consider their quantificational properties (in this case, ignoring the implicature). This idea is compatible with the Crain et al. plausible dissent hypothesis, in that it suggests that children exclude interpretations that are not compatible with any contextually relevant question. Thus, children's symmetrical interpretation of quantifiers may be due to a failure to converge on the speaker's intended QUD because this QUD seems implausible in the context, but also because additional information (e.g., extra elephants) is proffered by the experimenter, indicating its relevance to an alternative QUD (compatible with the Gricean maxim of relation). Second, the findings suggest that to test this idea in the case of quantifier spreading, it is necessary to present children with alternatives that are relevant to a quantificational QUD, and not to orthogonal non-quantificational questions like object kind. This is because when children evaluate quantificational expressions on the basis of object kind, they may ignore their quantificational properties, making it impossible to know whether or not their judgments reflect adult-like representations. Given these conclusions, we hypothesize that children entertain only plausible QUDs, compatible with Crain et al., but also that to test this question, children must be presented with utterances that satisfy plausible dissent in a way that is germane to the experimenter's intended QUD, an approach which has not been taken in previous studies.

In the present study, we addressed this idea by presenting children with utterances that satisfied plausible dissent, but that did so in a way that supported QUDs relevant to quantification, rather than object kinds. This allowed us to test the hypothesis that children's errors are remedied when they comprehend (and respond to) the experimenter's intended question, and therefore whether the problem of quantifier spreading is indeed a purely pragmatic

phenomenon. First, in Experiment 1, we tested the prediction that providing circumstances supporting the quantificational QUD (i.e., Do all girls ride an elephant?) might allow children to accept universally quantified statements. To do this, we tested children with scenes like the one in Figure 2 using sentences like “Every girl is riding an elephant”, but prior to test trials we introduced children to scenes that we called “falsifying alternatives”, in which there were, e.g., three girls, four elephants, and only two of these girls were riding elephants, as in Figure 3. We hypothesized that these scenes would both render the utterances felicitous - satisfying plausible dissent - and disambiguate the speaker’s intended quantificational QUD. Critically, on theories that propose a semantic deficit this manipulation should be unhelpful, since on these accounts, rejecting the sentences in falsifying alternative contexts should be compatible with children’s existing semantics, and therefore should not change how they interpret subsequent sentences on critical trials. However, on a QUD account, the availability of falsifying alternatives is informative, because it indicates the nature of the QUD – that it pertains to whether these three specific girls are on elephants or off of them – not whether elephants should have additional riders, or whether other animals should be present, etc. This, in turn, should allow children to interpret the statement relative to the falsifying alternative, rather than against other imagined alternatives (e.g., that might be otherwise triggered by the proffering of extra elephants, etc.).

In Experiment 2, we explored a separate way in which children’s failure to grasp the QUD might cause them to generate alternatives in a non-adult-like manner. As Drozd and van Loosbroek (2006) have suggested, children might assume upon seeing an extra elephant that the elephant is present because an additional (yet to be seen) girl is missing (see Aravind et al., 2017, for a similar idea). For example, upon hearing the sentence “Every girl is riding an elephant” in a context like that in Figure 2, they might reason that there exists some girl who owns the fourth

elephant who is not riding hers, and that the statement is therefore false. In Experiment 2, prior to viewing scenes like the one in Figure 2, children were presented with three characters in isolation, and told explicitly that these were the only three girls that existed, and that there are no others (what we called the *Domain Restriction* condition). We reasoned that if children judge universally quantified statements to be false because they posit the existence of a fourth, unseen, elephant-owning girl, then this tendency should decrease significantly when it is highlighted that only three girls exist (i.e., when the semantic domain is restricted to a world in which there are only three girls). Experiment 2 also included a condition that combined the falsifying alternatives and domain restriction manipulations to test whether they explained unique variance. Finally, we added a third condition to test whether the effect of providing falsifying alternatives at familiarization was specific to the specific objects mentioned – e.g., girls riding elephants – or whether the effect was more general, such that exposure to *any* scene that provided a basis for rejecting a universally quantified expression would also signal the intended QUD. Across these two experiments, we show that providing falsifying alternatives almost completely erases children’s difficulty with universally quantified sentences, and that any residual difference from adults is likely explained by difficulties with domain restriction.

Experiment 1

Experiment 1 consisted of three conditions. First, we sought to replicate the children’s classic symmetrical treatment of universal quantifiers, by asking them questions such as, “Is every girl riding an elephant?” in a context in which three girls are riding four elephants (as in Figure 2). We called this the *Baseline* condition. Second, we included a condition in which children received the same test trials as in the *Baseline*, but first were exposed to a familiarization scene in which, e.g., only two out of three girls were riding an elephant, in a

scene depicting four elephants. We called this the *Falsifying Alternatives* condition. Finally, a third condition also tested children as in the *Baseline* case, and like the *Falsifying Alternatives* condition included Familiarization trials that provided other alternative states of affairs; crucially, in this case the alternatives were uninformative with respect to identifying the question under discussion. Specifically, at familiarization, children were shown a scene in which, e.g., three out of three girls were riding three elephants, such that there were no leftovers. We called this the *Non-falsifying Alternatives* condition, since it provided no information regarding the nature of the question under discussion (i.e., whether the speaker was concerned with the types of animals, the number of girl riders, or whether contextually available girls were each riding elephants). The rationale of this third condition was to provide a control for the Falsifying Alternatives condition which, like it, included a familiarization phase. This was meant to ensure that any difference between the Falsifying Alternatives condition and Baseline performance was not due to a difference in familiarity with images or trial structure.

Method

Participants

We tested 146 typically developing preschoolers between the ages of 4 and 6 (4;0 – 6;0, $M = 5;0$), whose primary language was English. They were recruited from daycare centers in the San Diego area and the Reuben H. Fleet Science Center in San Diego. Based on criteria established before any examination of data, we excluded any children ($N=15$) who had either a systematic “No” bias or a systematic “Yes” bias across all trials, which we defined as providing an identical response across all trials. This resulted in the removal of 15 children. We also excluded children if they did not speak English as their primary language ($n = 8$), due to

experimenter error ($n = 2$), and due to ASD diagnosis ($n = 1$). The remaining children ($n = 120$) were divided into equal groups of 40, for each of the three conditions described below.

Stimuli and Procedure

Stimuli consisted of a series of 10 MS Power-point slides depicting clipart. Half of the slides depicted Critical trials in which three girl cartoon characters engaged in some sort of activity (riding, driving, etc.) each on her own ‘mount’ (e.g., elephant), or vehicle. A fourth mount or vehicle of the exact same kind but without a rider/driver was also present in each Critical trial scene. The other half of the slides depicted Control trials in which 3 or 4 animal characters, different for each trial, either featured a certain characteristic (e.g., having brown fur) or were engaged in an activity that was described by an intransitive verb (e.g., sleeping). Two of the Control trials depicted all of the characters having the characteristic / engaging in the activity and three of the trials depicted a proper subset of the characters having the characteristic / engaging in the activity. All trials were accompanied by a *Yes/No* question that asked children to evaluate whether a universally quantified proposition was true for the picture or not. For example, the corresponding question for the girls-riding-elephants slide (Critical trial) was “Is every girl riding an elephant?”; the corresponding question for the cats-sleeping slide (Control trial) was “Is every cat sleeping?” and so on. The stimuli were assembled to create two orders, one the reverse of the other. In both orders, trials alternated systematically between Critical and Control trials. Examples of the slides and corresponding questions can be seen in Figure 3 below.

Figure 3. Examples of Familiarization and Critical trials in Experiment 1.

Participants were randomly assigned to one of three between-subjects conditions, with the constraint that groups were matched according to age. The *Baseline* condition used the paradigm exactly as described above. The *Falsifying Alternatives* condition was identical to the *Baseline* condition except that it featured the addition of a falsifying alternative familiarization trial before each critical trial. On these familiarization trials, children were shown a picture in which only 2 out of 3 girls engaged in a particular action (for example, riding an elephant), but there were the same number of mounts or vehicles as before (i.e., four). Children were asked the

same question that accompanied the corresponding critical trial (e.g., “Is every girl riding an elephant?”). This trial was intended to offer children an alternative state of affairs under which the test sentence is false, thereby allowing them to determine that the question under discussion related to the number of girls riding elephants. Finally, the *Non-falsifying Alternatives* condition was identical to the *Falsifying Alternatives* condition except that the Familiarization trial occurring before the critical trial featured a “non-falsifying alternative”. The non-falsifying alternative consisted of a symmetrical picture where 3 out of 3 girls engaged in a particular action (for example, riding an elephant), but no additional mount or vehicle was present. This was intended to offer children an alternative state of affairs with respect to which the universally quantified utterance was clearly true, but nevertheless not useful for establishing plausible dissent. It also allowed us to ask whether mere familiarity with the paradigm or the quantified propositions themselves, rather than the specific nature of alternatives, might explain any improvements we witness. Tables featuring the entire content and structure of all experimental conditions can be found in Appendix A.

Children sat in front of a laptop computer in a quiet room in their preschool, or in a quiet area at the science museum. An experimenter sitting next to them administered the experiment by showing children the slides on the computer screen, asking them the corresponding questions from a prepared script, and recording children’s responses on an answer sheet. Children were told that they would play a “game” on the computer. The experimenter explained to them that they would look at some pictures together and that she was going to ask them some questions about the pictures; she explained that some of the questions might seem a little silly, but they should just answer what they thought. Children were shown each picture and the experimenter asked the corresponding *Yes/No* question, for instance, “Is every girl riding an elephant?”. In the

case of a *No* response, the experimenter also asked the child to explain their judgment. When the child provided a rationale for their response, it was recorded by the experimenter.

Results

Using separate Generalized Linear Mixed Models analyses in R, we predicted the frequency of correct and incorrect responses for control and critical trials, with Condition (*Baseline, Falsifying Alternatives, Non-falsifying Alternatives*) as a fixed factor and Subject, and Trial (Item) as random factors. Children performed similarly across conditions in Control trials ($M_{\text{BASELINE}} = .92$, $M_{\text{FALSIFYING ALTERNATIVES}} = .93$, $M_{\text{NON-FALSIFYING ALTERNATIVES}} = .95$; $\chi^2(2) = 0.889$, $p = .641$). Their performance in Critical trials, however, differed across conditions with the analysis revealing a main effect of Condition ($\chi^2(2) = 81.811$, $p < .0001$), driven by better performance in the *Falsifying Alternatives* condition ($M = .78$) relative to the *Baseline* ($M = .33$; $p < .0001$) and *Non-falsifying Alternatives* ($M = .12$; $p < .0001$) conditions. There was no significant difference between the *Baseline* and the *Non-falsifying Alternatives* conditions ($p = .354$). Figure 4 summarizes the results averaged across subjects and trials.

Figure 4. Performance of children in Experiment 1, as proportion of correct responses across conditions and trial types. Error bars represent standard errors.

Discussion

In Experiment 1, we replicated previously reported difficulties with quantifier spreading in our *Baseline* condition, but also found that these difficulties were almost completely – though not entirely – erased with modest changes to the paradigm. In particular, we found that providing states of affairs under which the utterances might be judged false significantly reduced the likelihood that children would later judge them to be false in the presence of an extra mount or vehicle. In contrast, providing children with non-falsifying alternatives – i.e., another situation in which the utterance is true – had no effect on children (and if anything made them worse than in

the *Baseline*).³ Overall, our results are compatible with the hypothesis that quantifier spreading is substantially reduced when the experimenter's intended QUD is made explicit by a preceding context that contrasts quantificationally from the test sentences. Still, although performance in the *Falsifying Alternatives* condition was strong, it was not perfect, raising the possibility that other factors may also contribute to children's difficulties. As noted in the Introduction, Drozd and van Loosbroek (2006) suggest that, upon seeing an extra elephant, children might assume that it is present because an additional (yet to be seen) girl is missing. In such a scenario, the statement "Every girl is riding an elephant" would indeed be false, making children's judgment entirely felicitous. Critically, children who are predisposed to infer additional girls might not be less likely to do so after viewing our falsifying alternatives trials, since such trials are false regardless of whether there are extra riders. Consequently, the remaining difference we see between the *Falsifying* and *Non-falsifying Alternatives* conditions may be explained by some small number of such children.

To investigate this possibility, Experiment 2 tested children with what we called a *Domain Restriction* condition, in which we altered the script to tell children explicitly that there were only three girls in the domain of discourse. To do this, we showed them an introductory image prior to the test trials in which there were three characters and told them that these were the only such characters in this context. In addition, we reminded them at the beginning of each critical trial that we would be talking about the three girls introduced at the start of the trial. We reasoned that if children spontaneously invoke a fourth character in the *Baseline* condition, this

³ Although this result was not significant, based on previous studies we know that the salience of an extra object may be increased by adding more extra objects (see General Discussion). By extension, it is possible that the contrast between the presence of an extra object at test vs. no extra object in familiarization could increase the likelihood of children's errors.

should be reduced in the *Domain Restriction* condition, leading to fewer “no” responses. In addition to this, we explored the extent to which our domain restriction manipulation had an independent effect relative to the *Falsifying Alternatives* condition of Experiment 1. One possibility, for example, is that the *Falsifying Alternatives* condition not only disambiguated the QUD, but also reinforced to children that only three girls were relevant to assessing the utterances. To test whether these factors were independent, we created a hybrid condition that we called *Domain Restriction + Falsifying Alternatives* in which both falsifying alternatives (the manipulation from the *Falsifying Alternatives* condition of Experiment 1) and domain restriction were provided before critical trials. This allowed us to ask whether the two manipulations, when combined, had a greater effect than either in isolation, and thus whether they contributed separate variance.

Finally, in addition to asking these questions, Experiment 2 asked whether the effect of providing falsifying alternatives at familiarization was specific to the particular items mentioned – e.g., girls riding elephants – or whether instead it engendered a more general shift in children’s reasoning, such that exposure to *any* scene that provided a basis for rejecting a universally quantified expression might help children identify the speaker’s intended QUD (e.g., such that exposure to scenes which render false “Every girl is riding an elephant” facilitates generation of scenes in which, e.g., it’s not the case that “Every *p* is *x*-ing a *q*”).

To probe this, we replicated the combined *Domain Restriction + Falsifying Alternatives* condition in a new condition that pseudo-randomized the order of familiarization trials, such that particular items mentioned at familiarization did not match the items mentioned at test on critical trials in the sense that the particular falsifying alternative preceding each critical trial did not match the following critical trial thematically. For instance, in the *Domain Restriction +*

Falsifying Alternatives condition of Experiment 1, critical trials that featured girls driving cars were always preceded by Familiarization trials that also featured girls driving cars. In contrast, in the *Domain Restriction + Scrambled Alternatives* condition of Experiment 2, critical trials that featured girls driving cars were preceded by Familiarization trials involving, for example, girls riding elephants.

We reasoned that if the effect of falsifying alternatives is mediated by reasoning exclusively about specific alternatives (e.g., girls riding elephants), then children should be less likely to answer ‘Yes’ in the critical trials of this condition. However, if the effect of the alternatives is not item-specific, but due to a more general shift in how children reason about the QUD of universally quantified utterances, then their performance should be equally good in the new condition.

Experiment 2

Our second experiment addressed three questions raised by Experiment 1. First, we asked whether domain restriction might be a factor behind children’s difficulties. Second, we asked whether any effects of domain restriction and alternatives might be independent – and thus might fully explain children’s behaviors when manipulated simultaneously – or whether both effects might have been captured in the *Falsifying Alternatives* condition. Finally, we asked whether children’s improvement in the *Falsifying Alternatives* condition of Experiment 1 was due to the specific alternatives provided or because of a more general shift in how children reason about the QUD in these universally quantified sentences.

Method

Participants

We tested 128 typically developing preschoolers between the ages of 4 and 6 (4;0 – 6;0, $M = 5;0$), whose primary language was English. As in Experiment 1, they were recruited from daycare centers in the San Diego area and the Reuben H. Fleet Science Center in San Diego. Based on criteria established before any examination of data, we excluded any children who displayed a systematic *Yes* or *No* bias, defined as children answering consistently *Yes* or *No* to every trial. This resulted in the removal of 4 children. We also excluded children whose primary language was not reported as English ($n = 2$), due to experimenter error ($n = 1$), or failure to complete the experiment ($n = 1$). The remaining children ($n = 120$) were divided into equal groups of 40, for each of the three conditions described below.

Stimuli & Procedure

We created stimuli for three between-subjects conditions. Stimuli for the *Domain Restriction* condition were identical to those in the *Baseline* condition except that they featured one additional slide at the beginning, which introduced the three girl characters as Mary, Jane and Lizzy, and explained that the game would involve seeing some pictures about what the girls were doing that day. Additionally, the critical trial questions were preceded by a short sentence reminding children that the question was about these three girls. For example, “Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?”. This manipulation was intended to make it clear to children that the three girls they saw on the pictures were the only girls relevant to the experiment, to help them restrict the domain of discourse to these three girls in exclusion to any others. Stimuli for the *Domain Restriction + Falsifying Alternatives* condition featured a combination of the domain restriction and falsifying alternatives manipulations of Experiment 1. In particular, the *Falsifying Alternatives* condition was modified to include the introduction and the domain-restricting prompts featured in the *Domain*

Restriction condition before each falsifying alternative and Critical trial. Full details on the change to the prompts can be found in Appendix A. Finally, stimuli for the *Domain Restriction + Scrambled Alternatives* condition modified the *Domain Restriction + Falsifying Alternatives* paradigm by scrambling the order of the falsifying alternatives so that they did not match the following critical trials. For example, the falsifying alternative with the girls riding elephants was followed by a critical trial with the girls driving cars.

Results

To assess the effects of the domain restriction manipulation, we analyzed the data by first comparing children's performance in the *Domain Restriction* condition with their performance in the *Baseline* condition of Experiment 1. Using separate Generalized Linear Mixed Models analyses in R, we predicted the frequency of correct and incorrect responses for control and critical trials, with Condition (*Baseline, Domain Restriction*) as a fixed factor and Subject and Item as random intercepts. Children's performance did not differ significantly between conditions for either the control ($M_{\text{BASELINE}} = .92$, $M_{\text{DOMAIN RESTRICTION}} = .99$; $\chi^2(1) = 2.608$, $p = 0.106$), or the critical trials ($M_{\text{BASELINE}} = .33$, $M_{\text{DOMAIN RESTRICTION}} = .49$; $\chi^2(1) = 0.395$, $p = 0.53$). This suggests that if domain restriction plays a role in this paradigm, its effect is relatively modest.

We next examined whether children's performance differed across the *Falsifying Alternatives* condition of Experiment 1 and the *Domain Restriction + Falsifying Alternatives* condition, in order to determine whether the combination of these two factors might have impacted children's judgments above and beyond falsifying alternatives alone. To do this, we used separate Generalized Linear Mixed Models analysis in R to predict the frequency of correct and incorrect responses for control and critical trials, with Condition (*Falsifying Alternatives,*

Domain Restriction + Falsifying Alternatives) as a fixed factor and Subject and Item as random intercepts. The analyses revealed that the performance of children on Control trials did not differ between conditions ($M_{\text{FALSIFYING ALTERNATIVES}} = .93$; $M_{\text{DOMAIN RESTRICTION + FALSIFYING ALTERNATIVES}} = .96$; $\chi^2(1) = 0.103, p = 0.748$). Likewise, performance on the Critical trials showed no difference ($M_{\text{FALSIFYING ALTERNATIVES}} = .78$; $M_{\text{DOMAIN RESTRICTION + FALSIFYING ALTERNATIVES}} = .92$; $\chi^2(1) = 0.706, p = 0.401$), although children were notably very near ceiling performance in this new experiment, unlike in Experiment 1. This again suggests that if domain restriction plays a role in children's judgments it is likely to be relatively small, and that the availability of falsifying alternatives is a more important factor.

Finally, we tested the possibility that falsifying alternatives had their effect because children reasoned exclusively about the specific alternatives provided, by comparing data from the *Domain Restriction + Falsifying Alternatives* condition to data from the *Domain Restriction + Scrambled Alternatives* condition. We again constructed separate models predicting the frequency of correct and incorrect responses for control and critical trials, with Condition (*Domain Restriction + Falsifying Alternatives, Domain Restriction + Scrambled Alternatives*) as a fixed factor and Subject and Item as random intercepts. The analyses revealed that children's performance was not significantly different across the two conditions in both the Control ($M_{\text{DOMAIN RESTRICTION + FALSIFYING ALTERNATIVES}} = .96$; $M_{\text{DOMAIN RESTRICTION + SCRAMBLED ALTERNATIVES}} = .92$; $\chi^2(1) = 0.053, p = 0.818$) and the Critical trials ($M_{\text{DOMAIN RESTRICTION + FALSIFYING ALTERNATIVES}} = .96$; $M_{\text{DOMAIN RESTRICTION + SCRAMBLED ALTERNATIVES}} = .92$; $\chi^2(1) = 0.043, p = 0.836$). These results suggest that scrambling falsifying alternatives at familiarization does not reduce their effect at test on critical trials, suggesting that their facilitating effect is due to a more general shift in how children reason about the QUD (which we discuss at further length in

the General Discussion). Figure 5 summarizes the results of Experiment 2 averaged across subjects and trials.

Figure 5. Performance of children in Experiment 2, as proportion of correct responses across conditions and trial types. Error bars represent standard errors.

Discussion

Experiment 2 found three main results. First, we found that emphasizing that only three characters were involved in the scenes did not reliably improve children's performance, suggesting that errors in the *Baseline* condition of Experiment 1 were likely not driven by problems of domain restriction (and the inference that extra girls must exist, that were not riding their assigned elephants). Relatedly, we found that combining the factors of domain restriction and falsifying alternatives into a new condition (*Domain Restriction + Falsifying Alternatives*) did not reliably improve children's performance compared to the independent manipulation of

the falsifying alternatives (*Falsifying Alternatives* condition – Experiment 1). Finally, scrambling the falsifying alternatives at familiarization did not affect children’s performance at test on critical trials, suggesting that the facilitating effect of the alternative states of affairs is not specific to the particular items presented, but engenders a more general effect on how children reason about the QUD of universally quantified statements.

General Discussion

We investigated children’s symmetrical interpretation of universal quantifiers – i.e., quantifier spreading – as a window into children’s comprehension of natural language quantification. In two studies, we tested the idea that these behaviors, first described by Piaget (1964), stem from the use of infelicitous questions that proffer irrelevant information, and thereby prevent children from converging on the experimenter’s intended “Question Under Discussion” (QUD). We asked whether these difficulties might be remedied by providing children with alternative states of affairs that disambiguate the speaker’s QUD. Also, we tested the narrower hypothesis that children’s difficulties stem from issues related to domain restriction, by explicitly telling children which entities to include in the domain of quantification.

First, in Experiment 1, we tested children with a Baseline condition that replicated the previously reported quantifier spreading phenomenon. Second, we tested children with either falsifying alternatives (i.e., states of affairs that rendered statements clearly false), or non-falsifying alternatives (that provided equivalent exposure to sentences and stimuli as the falsifying alternatives scenes, without providing falsifying alternative states of affairs). We found that children familiarized to falsifying alternatives scenes were subsequently much more likely to provide adult-like judgments on critical quantifier spreading trials at test, whereas familiarization to non-falsifying alternatives had no significant effect. Third, in Experiment 2, we showed that

domain restriction did not have a significant effect on quantifier spreading trials, though our condition that both restricted the domain and provided falsifying alternatives resulted in near ceiling performance on the task. We conclude from our domain restriction manipulations that our study does not entirely rule out some role for domain restriction, but that if such an effect exists it is relatively small, and much less robust than the effect of falsifying alternatives. Also, we showed that the facilitating effect of falsifying alternatives was not restricted to specific images, but that the effect was more general and abstract: falsifying alternatives were just as impactful to children when the types of objects depicted differed from those presented at test as when they matched. Further, this manipulation provided an additional replication of the effect of falsifying alternatives, for a total of three such demonstrations.

In the Introduction, we discussed the idea of Crain and colleagues (1996) that children's difficulties arise when experimental tasks violate "plausible dissent" – i.e., when children are asked a question which does not plausibly permit either a Yes or a No response. For example, in the context of our experiments, a question like, "Is every girl riding an elephant?" is infelicitous on Crain et al.'s analysis given that the experimenter can clearly see for themselves that every girl is riding an elephant. For Crain et al., the uncertainty created by this infelicity leads children to pragmatically infer an alternative basis for evaluating the question, leading to their non-adult-like behaviors. While we endorsed the intuition of Crain and colleagues' that children likely struggle to assess statements that are obviously true and violate plausible dissent, we also noted that plausible dissent can be satisfied in ways that are actually not germane to assessing children's comprehension of quantifiers, or the experimenter's intended question. We therefore proposed a narrower analysis based on Roberts' notion of the Question Under Discussion (Roberts, 1996/2012). On this analysis, children seek to identify the speaker's intended QUD and

use both the speaker's statement and other proffered contextual information to do so. When the context makes available alternatives that are compatible with the speaker's intended QUD – like in our Falsifying Alternatives condition – children readily determine which dimension of the scene is relevant to evaluating the sentence (e.g., how many girls are riding elephants) and respond in an adult-like manner. However, when a self-evident question is asked and such alternatives are not accessible, or when irrelevant contextual information is proffered (e.g., extra elephants) children instead contemplate alternative QUDs, leading to a range of diverse interpretations (the significance of which we discuss below).

Critically, whereas our data can be explained by the QUD account, they are not readily explained by theories that appeal to problems originating in the grammar, or to processing difficulties. Neither the grammatical discontinuity accounts nor the processing accounts can easily explain the effect of non-linguistic alternative states of affairs described in our experiments – at least not without important supplementary hypotheses. For example, on the hypothesis that children have weak meanings for quantifiers like *every* (Drozd, 2001), both a weak and a strong interpretation of a proposition like “Every girl is riding an elephant” should be available. Consequently, falsifying alternatives should be compatible with children's existing semantics, and should therefore be incapable of changing the representations they deploy when making judgments on the task (i.e., because either a weak or a strong semantics should lead children to judge the statement as false when one girl is *not* riding an elephant). Similarly, processing accounts cannot easily explain why our Falsifying Alternatives condition, but not our Non-falsifying Alternatives condition, should result in different judgments by children across these cases, barring additional assumptions. Very generally, in order to explain our experimental effects, the representations proposed by an adequate theory must be incompatible with falsifying

alternatives in order to explain why these alternatives cause children to re-analyze their interpretation of the task.

Is there a way to reconcile our findings with prior accounts? A prominent feature of most previous explanations of quantifier spreading is that they exhibit a ‘dual’ nature, and posit both processes of disruption and repair. While variously attributing the origin of children’s problems to different parts of grammar or processing difficulties, many previous accounts explain children’s final behaviors as resulting from context-driven repair strategies that are pragmatic in nature: When grammar fails to give children the correct interpretation of a question, they use pragmatic inference to arrive at another interpretation. Thus, one possible approach might be to argue that children’s difficulties do arise from grammar, but also that strong pragmatic cues related to the QUD lead children to revise their interpretations. The main problem with this idea, however, is that it requires positing two distinct developmental discontinuities rather than just one. On our account, children differ from adults because they struggle to identify the intended QUD and require contextual alternatives to do so; once they identify the intended QUD, they perform at near ceiling levels. However, any hybrid account deploying our proposal would require both this – which alone can explain our data – and also the unnecessary addition of grammatical or processing differences.

As already noted, our data fail to find strong support for the idea that children’s primary difficulty stems from problems with domain restriction. According to one previous account (Drozd & van Loosbroek, 2006), children’s limited processing capacities impact their ability to successfully infer the speaker’s intended domain of reference with the consequence that they infer, for example, that not all girls are riding elephants because there exist additional girls not presented in the scene. We show, however, that explicitly restricting the domain of reference by

telling children that no other girls exist had only a modest effect on their judgments, which alone does not explain substantial variance in our experiments, whereas providing falsifying alternatives does (see also Drozd & van Loosbroek, 2006, for similar evidence). Children's primary difficulties are not due to difficulties inferring the intended number of entities (e.g., elephants and girls) in a domain. Instead, they are due to misinterpreting the speaker's statement, and relating it to an unintended QUD.

One question not addressed by this study is why children respond in the particular way that they do in studies of quantifier spreading. Why is it that, rather than simply responding randomly, or refusing to answer, children instead deny that, e.g., every girl is riding an elephant? The literature suggests many ideas in response to this question. For example, perhaps children have a preference for scenes that exhibit one-to-one correspondence (Brooks & Sekerina, 2006), an idea which might also explain conjunctive interpretation of disjunctive sentences (Paris, 1973; Singh et al., 2016; Tieu et al., 2017; Huang & Crain, 2020; Skordos et al., 2020), or biconditional interpretation of conditional statements (Paris, 1973; Romain et al., 1983; Barrouillet, Gauffroy, & Lecas, 2008; Gauffroy & Barrouillet, 2009; Markovits et al., 2016). Alternatively, perhaps the kind or number of extra objects impacts children's judgments by shifting their relative salience (Freeman et al., 1982; Brooks & Sekerina, 2005; Gouro, Norita, Nakajima & Arijji, 2001; Kang, 2001; Minai et al., 2012; O'Grady et al., 2010; Sekerina & Sauermann, 2015; Sekerina et al., 2018; Sugisaki & Isobe, 2001). Our suggestion here is that while it may be possible to answer such questions for particular experimental contexts, efforts to generalize a formal account of how context impacts children's responses to infelicitous questions are unlikely to be fruitful, especially if the goal of such work is to understand core, semantic, representations. Studies that present children with irrelevant information (e.g., extra elephants) and show how they respond to

this may provide important evidence that children accommodate violations of the Gricean maxim of manner. However, because the range of irrelevant information that can be proffered is unlimited in principle, it seems unlikely that we should expect a general account of how children should respond to infelicitous questions, short of a complete theory of human abductive inference. For this reason, our suggestion is that researchers should abandon methods that explicitly hinge on infelicitous questions entirely, and seek out paradigms that pose unambiguous questions, and that avoid intentionally proffering irrelevant and misleading information that might trigger alternative, unintended QUDs.

To understand this suggestion, it's important to note that many of Piaget's tasks exhibited a similar structure, which in some cases have been abandoned for this very reason, replaced with methods that have made substantially more progress. Consider, for example, what is perhaps Piaget's most famous task, used for 50 years in studies of so-called "number conservation" (e.g., Piaget, 1965). In these studies, much like in tests of quantifier comprehension, children were asked a question with a self-evident answer, and were proffered irrelevant information in the context that they sought to accommodate, leading them to provide responses that we now know tell us little about their semantic representation of number. Specifically, children were initially shown two parallel lines of objects placed in one-to-one correspondence and asked if the two sets were equal or if one had more than the other. Having correctly judged that they were equal, children then watched as the experimenter stretched out one line, and then repeated the same question again - e.g., "Does this line have more, or does this line have more? Or are they the same?" Now, despite having answered that they were equal only moments earlier, children as old as 6 years of age judged that the longer line contained more. However, subsequent studies (McGarrigle & Donaldson, 1978; Mehler & Bever, 1967; Rose & Blank, 1974) showed that this

judgment was driven by an effort on the part of children to accommodate the experimenter's transformation of the set - an effort to find a meaning for "more" that made sense of why the experimenter might ask this question again, following the lengthening of one line. First, in a study by Rose and Blank (1974), when children were asked only the second question, after the transformation, they made far fewer errors, and were much more likely to judge that the two sets were equal, despite the fact that one set was longer than the other. Second, McGarrigle and Donaldson (1978) found that when the transformation occurred by accident, rather than intentionally, children again "conserved" far more frequently, judging that the two sets were equal despite being asked the same question twice in succession. Finally, in a study by Mehler and Bever (1967), when children as young as 3 years of age were shown lines of candy and simply asked which row they preferred, they succeeded robustly. Of significance to the present discussion, what these studies show is that conservation failures ultimately tell us very little about children's understanding of equality. Further, despite decades of studies, the precise reason that children, but not adults, make these errors remains poorly understood. Although we know that these errors arise from a pragmatic violation, we still have no consensus account of why children respond in the specific way they do, while studies using alternative methods have made considerable progress in understanding how children represent number (for a review of the literature on one-to-one correspondence beginning with Piaget, see Schneider et al., 2022). Very generally, this example suggests that efforts to provide a formal account of how children respond to infelicitous questions may not be fruitful, especially if the goal of such work is to understand core, semantic, representations, rather than how children accommodate strange questions.

Lessons from this study may also extend to other developmental phenomena regarding children's reasoning with quantified sentences. In particular, our account flows naturally from an

access-to-alternatives approach that has been fruitfully applied to case studies such as scalar implicature, negation, and conditional statements and negation (see Skordos & Barner, 2019, for discussion). According to this idea, children face difficulties when reasoning with quantified statements largely because of problems in restricting the inferential hypothesis space in language comprehension (Bale & Barner, 2013; Barner, Brooks, & Bale, 2011; Chierchia et al., 2001; Foppolo et al., 2012; Gotzner, Barner, & Crain, 2020; Hochstein et al., 2014; Singh et al., 2016; Skordos, Feiman, Bale, & Barner, 2021; Skordos & Papafragou, 2016; Tieu et al., 2015; 2017). To comprehend an utterance, the listener attempts to identify its meaning from a large space of possible alternative meanings. Although this hypothesis space is narrowed by the semantics of the propositions involved, the space of possibilities nevertheless remains large: Any individual utterance may be compatible with a host of distinct QUDs, each of which is compatible with different sets of alternatives that might be considered. For this reason, children’s difficulties with linguistic inferences like scalar implicature can be mitigated either by making available specific linguistic alternatives (which narrow the set of likely QUDs), or by directly indicating the intended QUD (allowing the child to spontaneously identify relevant linguistic alternatives).

- (3) a. Some of the animals are holding umbrellas
- b. Are any/no animals holding umbrellas?
- c. Are some/all animals holding umbrellas?
- d. Some of the animals are holding shovels.
- e. All of the animals are holding umbrellas.

For example, given a sentence like the one in (3a) (i.e., “Some animals are holding umbrellas”), where focus can be placed either on the quantificational element “some” or the kind term “umbrellas”, different QUDs may result as shown in (3b) and (3c). If children favor the

QUD “Are any/no animals holding umbrellas?” over “Are some/all animals holding umbrellas?” then they may accept the statement in (3a) even when all animals are holding umbrellas, resulting in the appearance of an SI failure. Whereas providing children with linguistic alternatives like in (3d) that replace “umbrella” with “shovel” has no impact on scalar implicature, providing them with an alternative like in (3e) that contains the word “all” leads children to compute implicatures at much higher rates. On our account, this is because providing such alternatives disambiguates the intended QUD (for experimental evidence, see Skordos & Papafragou, 2016).

Studies of other case studies also find, as we do here, that providing children with experimental contexts that support dissent – and in doing so expose the intended QUD – lead to more adult-like patterns of inference. These studies include work on scope ambiguity (Di Bacco et al., 2018; Gualmini, 2004; Gualmini et al., 2008; Musolino & Lidz, 2003), negation (Nordmeyer & Frank, 2018), and conditional statements (see Skordos & Barner, 2019, for discussion). In studies of negation (Nordmeyer & Frank, 2014; Reuter, Feiman & Snedeker, 2018), and conditional inferences (Rumain, Connell, & Braine, 1983), when children are presented with states of affairs that are compatible with alternative linguistic descriptions that disambiguate a speaker’s intended QUD, they subsequently perform significantly better than in baseline conditions. For example, Rumain et al. (1983) found that children presented with the conditional in (4) incorrectly inferred that it must have rained if told that John brought his umbrella, but did so significantly less if also presented with a statement like the one in (4b). Similarly, Reuter et al. (2018) found that although 2-year-old children incorrectly looked more to a broken plate than an unbroken one when presented with a negated statement like the one in (5a), they performed significantly better when this trial was preceded by one in which the

affirmative statement was true, and thus children had a semantic model compatible with the affirmative, against which they could evaluate the negated statement.

- (4) a. If it rains, John may bring his umbrella.
- b. If it doesn't rain, John may bring his umbrella.
- c. John brought his umbrella: Did it rain?
- (5) a. The girl didn't break one of the plates!
- b. The girl broke one of the plates.

In summary, previous studies suggest that a broad and parsimonious account that appeals to the QUD may be possible, and that it may not be necessary to posit idiosyncratic accounts for each of children's difficulties with logical reasoning – whether implicature, conditionals, negation, or scalar implicature. Future studies should explore how this approach can be extended to additional case studies, while also seeking more direct evidence for the idea that problems identifying the QUD might also explain children's difficulties with pragmatic inference in general.

Acknowledgements

This work was supported in part by an award from the James S. McDonnell Foundation to (last author). The authors would like to thank the Fleet Science Center in SD, as well as all preschools and families who participated.

References

- Aravind, A, De Villiers, J., de Villiers, P., Lonigan, C., Phillips, B., Clancy, J., Landry, S. et al. (2017). Children's quantification with every over time. *Glossa*, 10.
- Austin, K., Theakston, A., Lieven, E. & Tomasello, M. Young children's understanding of denial. *Dev. Psychol.* 50, 2061–2070 (2014).
- Bale, A., & Barner, D. (2013). Grammatical alternatives and pragmatic development. In *Alternatives in semantics* (pp. 238-266). Palgrave Macmillan, London.
- Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, 188, 87-96.
- Barner, D., Chow, K., & Yang, S., J. (2009). Finding one's meaning: A test of the relation between quantifiers and integers in language development. *Cognitive Psychology*, 58, 195–219.
- Barner, D., Hochstein, L. K., Rubenson, M. P., & Bale, A. (2018). Four-year-old children compute scalar implicatures in absence of epistemic reasoning. *Semantics in language acquisition*, 24, 325.
- Braine, M., & Romain, B. (1981). Children's comprehension of “or”: evidence for a sequence of competencies. *Journal of Experimental Child Psychology*, 31, 46-70.
- Brooks, P. J., & Parshina, O. (2019). Quantifier Spreading. In *The Oxford Handbook of Experimental Semantics and Pragmatics*.
- Brooks, P., J. & Sekerina, I. (2005). Shortcuts to quantifier interpretation in children and adults. *Language Acquisition*, 13, 177-206.

- Brunetti, M., Di Matteo, R., Aureli, T., Garito, M. C., & Casadio, C. (2021). Class inclusion versus quantifiers comprehension tasks: an experimental study with school-aged children. *Cognitive Processing*, 22(1), 151-158.
- Cesana-Arlotti, N., Kovács, Á. M., & Téglás, E. (2020). Infants recruit logic to learn about the social world. *Nature communications*, 11(1), 1-9.
- Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., & Bonatti, L. L. (2018). Precursors of logical reasoning in preverbal human infants. *Science*, 359(6381), 1263-1266.
- Chen, F., Rosenstein, L., & Hackl, M. (2021). Quantifier-Spreading under Negation.
- Chierchia, G., Crain, S., Guasti, M., T., Gualmini, A., & Meroni L. (2001). The acquisition of disjunction: Evidence for a grammatical view of scalar implicatures. *Proceedings from the Annual Boston University Conference on Language Development*, 25, 157-168.
- Clothier, K. K. (2019). The many meanings of many: A pragmatic-semantic account of ambiguous, vague quantifiers with experimental support (Doctoral dissertation, Johns Hopkins University).
- Cohen, A. (2001). Relative readings of many, often, and generics. *Natural Language Semantics* 9. 41-67.
- Crain, S., Thornton, R., Boster, C., Conway, L., Lillo-Martin, D., & Woodams, E. (1996). Quantification without qualification. *Language Acquisition*, 5, 83-153.
- Denić, M., & Chemla, E. (2020). Quantifier Spreading in Child Language as Distributive Inferences. *Linguistic Inquiry*, 51(1), 141-153.

- Di Bacco, F., Tieu, L., Moscati, V., Folli, R., Sevdali, C., & Romoli, J. (2017). Testing the QUD approach: Children's comprehension of scopally ambiguous questions. *In Proceedings of the 34th West Coast Conference on Formal Linguistics*, April 29-May 1, 2016, Salt Lake City, Utah, US (pp. 177-186).
- Donaldson, M., & McGarrigle, J. (1974). Some clues to the nature of semantic development. *Journal of Child Language*, 1(2), 185-194.
- Donaldson, M., & Lloyd, P. (1974). *Sentences and situations: children's judgments of match and mismatch*. Centre National de la Recherche Scientifique.
- Drozd, K. F. (2001). Children's weak interpretations of universally quantified questions. In M. Bowerman and S. C. Levinson (eds.), *Language acquisition and conceptual development* (pp. 340-376). Cambridge: Cambridge University Press.
- Drozd, K., Anđelković, D., Savić, M., Tošković, O., Gavarró, A., Lite, A., ... & Sauerland, U. (2019). A crosslinguistic study of symmetrical judgments. *Proceedings of the 43rd Boston University Conference on Language Development*, ed. Megan M. Brown and Brady Dailey, 217-230. Somerville, MA: Cascadilla Press.
- Drozd, K. & van Loosbroek, E. (1999). "Weak quantification, plausible dissent, and the development of children's pragmatic competence." In A. Greenhill, H. Littlefield, and Ch. Tano (Eds), *Proceedings of the 23d Annual Boston University Conference on Language Development*. (pp. 184-195) Somerville MA: Cascadilla Press.
- Drozd, K. F., & van Loosbroek, E. (2006). The effect of context on children's interpretations of universally quantified sentences. In: *Semantics in Acquisition* (pp. 115-140). Springer Netherlands

- Feeney, A., Scafton, S., Duckworth, A., & Handley, S. (2004). The story of *some*: everyday pragmatic inference by children and adults. *Canadian Journal of Experimental Psychology*, 58, 121-132.
- Feiman, R., Mody, S., Sanborn, S. & Carey, S. What do you mean, no? Toddlers' comprehension of logical "no" and "not". *Lang. Learn. Dev.* 13, 430–450 (2017).
- Foppolo, F., Guasti, M.T., & Chierchia, G. (2012). Scalar implicatures in child language: Give children a chance. *Language Learning and Development*, 8, 365-394.
- Freeman, N., H., Sinha, C., G., & Stedmon, A., J. (1982). All the cars— which cars? From word meaning to discourse meaning to discourse analysis. In M. Beveridge (ed.), *Children thinking through language* (pp. 52-74). London: Edward Arnold.
- Geurts, B. (2003). Quantifying kids. *Language Acquisition*, 11, 197-218.
- Gotzner, N., Barner, D., & Crain, S. (2020). Disjunction triggers exhaustivity implicatures in 4-to 5-year-olds: Investigating the role of access to alternatives. *Journal of Semantics*, 37(2), 219-245.
- Gualmini, A. (2004) The Ups and Downs of Child Language, *Outstanding Dissertations in Linguistics* series, Routledge, New York, NY.
- Gualmini, A., Meroni, L., & Crain, S. (2003). An asymmetric universal in child language. In *Proceedings of Sinn und Bedeutung* (Vol. 7, pp. 136-148).
- Gualmini, A., Hulsey, S., Hacquard, V., & Fox, D. (2008). The question–answer requirement for scope assignment. *Natural language semantics*, 16(3), 205-237.
- Guasti, M. T., Chierchia, G., Crain, S., Foppolo F., Gualmini, A., & Meroni L. (2005). Why children and adults sometimes (but not always) compute implicatures. *Language and Cognitive Processes*, 20, 667-696.

- Gouro, T., Norita, H., Nakajima, M., & Arijji, K. I. (2001). Children's interpretation of universal quantifier and pragmatic interference. In Proceedings of the Second Tokyo University Conference on Psycholinguistics (pp. 61-78).
- Hochstein, L., Bale, A., Fox, D., & Barner, D. (2016). Ignorance and inference: do problems with Gricean epistemic reasoning explain children's difficulty with scalar implicature?. *Journal of Semantics*, 33(1), 107-135.
- Hollebrandse, B. (2004). Topichood and quantification in L1 Dutch. *International Review of Applied Linguistics in Language Teaching* 42(2), 203–215.
- Horn, L. R. (1972). *On the semantic properties of the logical operators in English*. Doctoral diss., UCLA.
- Horowitz, A. C., Schneider, R. M., & Frank, M. C. (2018). The trouble with quantifiers: exploring children's deficits in scalar implicature. *Child development*, 89(6), e572-e593.
- Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child: Classification and seriation*. London: Routledge and Kegan Paul.
- Kang, H. K. (2001). Quantifier spreading: Linguistic and pragmatic considerations. *Lingua*, 111(8), 591-627.
- Katsos, N., & Bishop, D. V. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition*, 120, 67-81.
- Ke, A. H., & Gao, L. (2020). Domain restriction in child Mandarin: Implications for quantifier spreading. *Linguistics*, 58(6), 1839-1875.

Kiss, K., & Zétényi, T. (2017). Quantifier spreading: children misled by ostensive cues. *Glossa: a journal of general linguistics*, 2(1), 1-20.

Musolino, J., & Gualmini, A. (2004). The role of partitivity in child language. *Language Acquisition*, 12(1), 97-107.

Minai, U., Jincho, N., Yamane, N., & Mazuka, R. (2012). What hinders child semantic computation: children's universal quantification and the development of cognitive control. *Journal of Child Language*, 39(5), 919-956.

Mody, S. & Carey, S. The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition* 154, 40–48 (2016).

Morris, B. J. (2008). Logically speaking: evidence for item-based acquisition of the connectives AND & OR. *J. Cogn. Dev.* 9, 67–88.

Musolino, J., & Lidz, J. (2003). The scope of isomorphism: Turning adults into children. *Language Acquisition*, 11(4), 277-291.

Musolino, J. & Lidz, J. (2006). Why children aren't universally successful with quantification. *Linguistics* 44, 817–852.

Nordmeyer, A. E., & Frank, M. C. (2014). The role of context in young children's comprehension of negation. *Journal of Memory and Language*, 77, 25–39.

Nordmeyer, A. E. & Frank, M. C. (2014). The role of context in young children's comprehension of negation. *J. Mem. Lang.* 77, 25–39.

Nordmeyer, A. E., & Frank, M. C. (2018). Early Understanding of Pragmatic Principles in Children's Judgments of Negative Sentences. *Language Learning and Development*, 14(4), 262-278.

- Noveck, I. (2001). When children are more logical than adults: Experimental investigations of scalar implicatures. *Cognition*, 78, 165-188.
- O'Grady, W., Suzuki, T., & Yoshinaga, N. (2010). Quantifier spreading: New evidence from Japanese. *Language Learning and Development*, 6(2), 116-125.
- Ozturk, O., & Papafragou, A. (2015). The acquisition of epistemic modality: From semantic meaning to pragmatic interpretation. *Language Learning and Development*, 11, 3, 191-214.
- Papafragou, A., & Musolino, J. (2003). Scalar implicatures: Experiments at the semantics pragmatics interface. *Cognition*, 86, 253-282.
- Philip, W. (1995). *Event quantification in the acquisition of universal quantification*. Doctoral dissertation, University of Massachusetts–Amherst
- Philip, W. (2003). Specific indefinites and quantifier scope for children acquiring Dutch and Chinese. Paper presented at the 2002 Meeting of the Linguistic Society of the Netherlands.
- Philip, W., & Verrips, M. (1994). Dutch preschoolers' elke. In *Boston University Conference on Language Development*.
- Pouscoulous, N., Noveck, I., Politzer, G., Bastide, A. (2007). Processing costs and implicature development. *Language Acquisition*, 14, 347-375.
- Reinhart, T. (2004). The processing cost of reference-set computation: Acquisition of stress shift and focus. *Language Acquisition*, 12(2), 109–155.

- Reuter, T., Feiman, R., & Snedeker, J. (2018). Getting to No: Pragmatic and Semantic Factors in Two- and Three- Year-Olds' Understanding of Negation. *Child development*, 89(4), e364-e381.
- Roberts, C. (2012). Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5, 6-1
- Roeper, T., & De Villiers, J. (1993). The emergence of bound variable structures. In Knowledge and language (pp. 105-139). Springer Netherlands.. In E. Reuland and W. Abraham (Eds), *Knowledge and Language: from Orwell's Problem to Plato's Problem*. (pp. 105–139), Dordrecht: Kluwer.
- Roeper, T., Strauss, U., & Pearson, B. Z. (2006). The acquisition path of the determiner quantifier every: Two kinds of spreading. *Papers in Language Acquisition*, University of Massachusetts Occasional Papers UMOP, 34.
- Schneider, R. M., Brockbank, E., Feiman, R., & Barner, D. (2022). Counting and the ontogenetic origins of exact equality. *Cognition*, 218, 104952.
- Sekerina, I. A., Brooks, P. J., Campanelli, L., & Schwartz, A. M. (2018). Quantifier Spreading in School-Age Children: An Eye-Tracking Study. In *Linguistic and Cognitive Aspects of Quantification* (pp. 171-192). Springer, Cham.
- Sekerina, I. A., & Sauermann, A. (2017). Quantifier spreading in child eye movements: A case of the Russian quantifier *kazhdij* 'every'.

- Singh, R., Wexler, K., Astle-Rahim, A., Kamawar, D. & Fox, D. (2016). Children interpret disjunction as conjunction: Consequences for theories of implicature and child development. *Nat. Lang. Semant.* 24, 305–352.
- Skordos, D., & Barner, D. (2019) Language comprehension, inference, and alternatives. In C. Cummins & N. Katsos (Eds.), *Oxford Handbook of Experimental Semantics and Pragmatics*. Oxford: Oxford University Press.
- Skordos, D., Feiman, R., Bale, A., & Barner, D. (2020). Do Children Interpret ‘or’ Conjunctively?. *Journal of Semantics*, 37(2), 247-267.
- Skordos, D., & Papafragou, A. (2016). Children’s derivation of scalar implicatures: alternatives and relevance. *Cognition*, 153: 6-18.
- Smith, C., L. (1980). Quantifiers and question answering in young children. *Journal of Experimental Child Psychology*, 30, 191–205.
- Stiller, A., Goodman, N. D., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language, Learning and Development*, 11, 2, 176-190.
- Sugisaki, K., & Isobe, M. (2001). Quantification without qualification without plausible dissent. *University of Massachusetts Occasional Papers in Linguistics*, 25, 97-100.
- Tieu, L., Romoli, J., Zhou, P., & Crain, S. (2015). Children's Knowledge of Free Choice Inferences and Scalar Implicatures. *Journal of Semantics*, 0, 1-30
- Tieu, L., Yatsushiro, K., Cremers, A., Romoli, J., Sauerland, U., & Chemla, E. (2017). On the role of alternatives in the acquisition of simple and complex disjunctions in French and Japanese. *Journal of Semantics*, 34(1), 127-152.

Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem and A. ter Meulen (Eds.), *Generalized Quantifiers in Natural Language*. (pp. 45-71) Dordrecht, The Netherlands: Foris Publications.

APPENDIX A

List of trial questions used for Experiments 1 and 2. The lists maintain the actual ordering used for the “Forward” order of the experiments and were reversed for the “Reverse” order (the ‘Alternative’ / ‘Critical’ trial pairings were maintained in the “Reverse” order).

Experiment 1

Baseline

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Critical	Is every girl riding an elephant?	Yes
Control	Is every cat sleeping?	No
Critical	Is every girl driving a car?	Yes
Control	Is every duck wearing a hat?	Yes
Critical	Is every girl climbing a tree?	Yes
Control	Does every dog have brown fur?	No
Critical	Is every girl sailing a boat?	Yes
Control	Is every fish orange?	Yes
Critical	Is every girl flying a plane?	Yes
Control	Is every turtle swimming in the water?	No

Falsifying Alternatives

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Alternative	Is every girl riding an elephant?	No
Critical	Is every girl riding an elephant?	Yes
Control	Is every cat sleeping?	No
Alternative	Is every girl driving a car?	No
Critical	Is every girl driving a car?	Yes
Control	Is every duck wearing a hat?	Yes
Alternative	Is every girl climbing a tree?	No
Critical	Is every girl climbing a tree?	Yes
Control	Does every dog have brown fur?	No
Alternative	Is every girl sailing a boat?	No
Critical	Is every girl sailing a boat?	Yes
Control	Is every fish orange?	Yes
Alternative	Is every girl flying a plane?	No
Critical	Is every girl flying a plane?	Yes
Control	Is every turtle swimming in the water?	No

Non-falsifying Alternatives

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Alternative	Is every girl riding an elephant?	Yes
Critical	Is every girl riding an elephant?	Yes
Control	Is every cat sleeping?	No
Alternative	Is every girl driving a car?	Yes
Critical	Is every girl driving a car?	Yes
Control	Is every duck wearing a hat?	Yes
Alternative	Is every girl climbing a tree?	Yes
Critical	Is every girl climbing a tree?	Yes
Control	Does every dog have brown fur?	No
Alternative	Is every girl sailing a boat?	Yes
Critical	Is every girl sailing a boat?	Yes
Control	Is every fish orange?	Yes
Alternative	Is every girl flying a plane?	Yes
Critical	Is every girl flying a plane?	Yes
Control	Is every turtle swimming in the water?	No

Experiment 2

Domain Restriction

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Critical	Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?	Yes
Control	These are Mary, Jane and Lizzy's cats. Is every cat sleeping?	No
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl driving a car?	Yes
Control	This is a pond in the girls' neighborhood. Is every duck wearing a hat?	Yes
Critical	Mary, Jane and Lizzy are playing at the park! Is every girl climbing a tree?	Yes
Control	These are Mary, Jane, and Lizzy's dogs. Does every dog have brown fur?	No
Critical	Now Mary, Jane and Lizzy are at the ocean! Is every girl sailing a boat?	Yes
Control	Mary, Jane and Lizzy saw some fish! Is every fish orange?	Yes
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl flying a plane?	Yes
Control	Look, there are some turtles at the beach where Mary, Jane and Lizzy live! Is every turtle swimming in the water?	No

Domain Restriction + Falsifying Alternatives

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Alternative	Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?	No
Critical	Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?	Yes
Control	These are Mary, Jane and Lizzy's cats. Is every cat sleeping?	No
Alternative	Let's see what Mary, Jane and Lizzy are doing now. Is every girl driving a car?	No
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl driving a car?	Yes
Control	This is a pond in the girls' neighborhood. Is every duck wearing a hat?	Yes
Alternative	Mary, Jane and Lizzy are playing at the park! Is every girl climbing a tree?	No
Critical	Mary, Jane and Lizzy are playing at the park! Is every girl climbing a tree?	Yes
Control	These are Mary, Jane, and Lizzy's dogs. Does every dog have brown fur?	No
Alternative	Now Mary, Jane and Lizzy are at the ocean! Is every girl sailing a boat?	No
Critical	Now Mary, Jane and Lizzy are at the ocean! Is every girl sailing a boat?	Yes
Control	Mary, Jane and Lizzy saw some fish! Is every fish orange?	Yes
Alternative	Let's see what Mary, Jane and Lizzy are doing now. Is every girl flying a plane?	No
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl flying a plane?	Yes
Control	Look, there are some turtles at the beach where Mary, Jane and Lizzy live! Is every turtle swimming in the water?	No

Domain Restriction + Scrambled Alternatives

<i>Trial</i>	<i>Prompt</i>	<i>Target Answer</i>
Alternative	Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?	No
Critical	Mary, Jane and Lizzy are playing at the park! Is every girl climbing a tree?	Yes
Control	These are Mary, Jane and Lizzy's cats. Is every cat sleeping?	No
Alternative	Let's see what Mary, Jane and Lizzy are doing now. Is every girl driving a car?	No
Critical	Mary, Jane and Lizzy are at the ocean! Is every girl sailing a boat?	Yes
Control	This is a pond in the girls' neighborhood. Is every duck wearing a hat?	Yes
Alternative	Mary, Jane and Lizzy are playing at the park! Is every girl climbing a tree?	No
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl flying a plane?	Yes

Control	These are Mary, Jane, and Lizzy's dogs. Does every dog have brown fur?	No
Alternative	Now Mary, Jane and Lizzy are at the ocean! Is every girl sailing a boat?	No
Critical	Look, Mary, Jane and Lizzy found some elephants! Is every girl riding an elephant?	Yes
Control	Mary, Jane and Lizzy saw some fish! Is every fish orange?	Yes
Alternative	Let's see what Mary, Jane and Lizzy are doing now. Is every girl flying a plane?	No
Critical	Let's see what Mary, Jane and Lizzy are doing now. Is every girl driving a car?	Yes
Control	Look, there are some turtles at the beach where Mary, Jane and Lizzy live! Is every turtle in the water?	No

APPENDIX B

Example of Falsifying Alternative paired with its corresponding Critical Trial

Is every girl riding an elephant?

Is every girl riding an elephant?