A continuity in logical development: domain-general logical inference by 2.5-year-old toddlers

Authors: xxxx

Affiliations: xxxx

Abstract

Adults can learn, form predictions and make decisions by applying domain-general logical inferences. The ontogenetic origins of these abilities remain unknown. Here we find that 2.5-year-old toddlers (N=36) can reason using Disjunctive Syllogism (i.e., A OR B, NOT A, THEREFORE B) across three contexts, which argues that domain-general logical reasoning may be in place from as early as the third year of life.

Keywords

Cognitive Development, Disjunctive Syllogism, Word Learning, Social Learning, Negation,

Statement of Relevance

Children grow up surrounded by opportunities to learn (the language of their community, the movements of their body, other people's preferences and mental lives, games, social norms, etc.). Here, we find that 2.5-year-old toddlers rely on a logical reasoning strategy, Disjunctive Syllogism, across a variety of situations, all before they have any formal education or extensive experience with words for expressing logical meanings. In learning new words, learning new facts about a person, and finding the winner of a race, toddlers systematically consider and reject competitors before deciding who must be the winner. This suggests that toddlers may have a general-purpose logical reasoning tool that they can use in any situation. Future work will determine what other logical tools toddlers can use.

Introduction

The capacity to think using abstract logical operations is at the core of some of our species' most powerful abilities. By representing logical relations between hypotheses, we can think not just about the way things are, but also about the ways things cannot be, and draw inferences about the way they must be. This capacity is seminal for both professional scientific reasoning and reasoning in daily life. We do this, for example, when we use Disjunctive Syllogism (e.g., EITHER the sun OR the earth is at the center of the planets' orbits; it is NOT the earth, because this would require implausibly complicated orbits; THEREFORE, the sun is at the center of the planets' orbits).

Even more powerfully, the human aptitude for domain-general learning, planning, and decision-making rests on a foundation of logical representations that allow us to reason about observations from any domain, from physics to knitting, using the same domain-general inference schemas (Fodor, 1979; Tarski & Corcoran, 1986). Where does this logical ability come from—is it a hard-won accomplishment of educated adults, or is it within the grasp even of young children? A handful of studies have asked whether children as young as 2.5 years can deploy the foundational logical schema Disjunctive Syllogism, but these studies have only tested its use within a single domain— the tracking of hidden

objects, and so cannot address the issue of domain generality (Gautam et al., 2021; Grigoroglou et al., 2019; Mody & Carey, 2016).

In a complementary line of work, Cesana-Arlotti and colleagues explored preverbal disjunctive reasoning in infants with tasks involving ambiguity about object identity (Cesana-Arlotti et al., 2018)). Crucially, in one study, 14-month-old infants learned the preference of an agent reaching for a goal-object whose identity was concealed - it had to be inferred using a Disjunctive Inference¹. (Cesana-Arlotti et al., 2020)). Infants success in this task can be explained with a from logical inference about an agent's preference (i.e., the agents prefers the toy car OR she prefer the ball; she avoided the ball; THEREFORE, she prefers the toy car). Alternatively, infants may first make a logical inference about the identity of the hidden object (i.e., the hidden object is the toy car OR the ball; the hidden object is a different object than the ball; THEREFORE, the hidden object is the toy car) and only afterward attribute the preference. Hence, this finding hints at a possible domain generality of infants' logical representation, but, clearly, targeted testing of domain generality is still required.

What would it mean for children to have a domain general understanding of Disjunctive Syllogism? At least, with adults, it is transparent that we know the words *or* and *not* and we can apply them to any domain – so we exhibit domain generality. But, the link between nonverbal Disjunctive Syllogism and children's understanding of verbal negation is not well understood. While verbal negation is powerful because it can apply to any domain (“not the ball,” “not Mommy,” “not loud”), it is unknown whether children's early understanding of verbal negation in fact implicates domain generality. Toddlers respond to sentences that explicitly deny the presence of an object at a location (i.e., “the ball is not in the bucket”) by avoiding the location and searching elsewhere (Austin et al., 2014; Feiman et al., 2017)). But this might reflect avoidance based on domain-specific representations (e.g., *avoid that location* or *that location is empty*), rather than the operation of elimination that can support a logical inference. Thus, it remains unclear whether early comprehension of verbal negation reflects logical understanding, and whether this understanding is related to that exhibited in nonverbal situations.

To find out, we asked whether toddlers could perform a Disjunctive Syllogism in three tasks that were identically structured except for the content toddlers were required to reason over, which varied widely (see supplementary materials). The Word-Learning task asked toddlers to identify familiar vs. unfamiliar objects (e.g., “Point at the ball/ dax” (Markman & Wachtel, 1988)) (n = 12; age range 2.3 – 2.8 years); the Social-Learning task required toddlers to identify animals based on their previously learned (i.e., familiar) versus novel (unfamiliar) preferences (e.g., “I like taking walks”) (n = 12; age range 2.5 – 3.4 years); the Explicit Negation task had toddlers identify the winner of a race using a linguistic assertion versus negation (“The winner is/ is not the bird”) (n = 12; age range 2.4 – 3.2 years). The set-up for each game was similar, and quite simple. In each task, across many trials, toddlers saw two options and had to choose just one of them. For example, in the Word Learning task, toddlers might see a ball and an unfamiliar object, and hear “Point at the dax!” If

¹ In a Disjunctive Inference, two alternatives are represented in a disjunctive relation, so one is inferred when the other is eliminated. The elimination of one option can be performed via mutually-exclusive representations (e.g., “it is a different object than the ball” rules out “it is the ball”). When the elimination is attained with logical negation (e.g., “it is NOT the ball” rules out “it is the ball”), we refer to the inference as (an instance of) the Disjunctive Syllogism. (see Feiman, Mody, and Carey, 2022, for a helpful discussion of a related distinction).

toddlers infer the right choice by rejecting the wrong one, they have performed a Disjunctive Syllogism (e.g., “dax” must refer to one object or the other; “dax” does not refer to the ball; therefore, “dax” must refer to the unfamiliar object).

Performing the Disjunctive Syllogism is not, however, the only way to succeed in our tasks. For example, toddlers might succeed on unfamiliar trials in the three tasks without using a Disjunctive Syllogism, simply by positively seeking an object that contrasts with the competitor – e.g., they need only “map novelty to novelty” (a positive strategy that does not predict double-checking the familiar object; (Mervis et al., 1994)). How can we ascertain whether toddlers use a logical inference in our tasks?

We predicted that if toddlers actively reason through the rejection of the competitor object (e.g., “dax” does not refer to the ball), as adults and older children do, then they may show an increased tendency to shift their gaze to the competitor as they think through this rejection (Halberda, 2006). We recorded toddlers’ eye movements throughout each trial. Inspired by work in adults and preschoolers (Halberda, 2006), we formalized the following prediction: if toddlers were already fixating the unfamiliar target object at the moment of hearing the unfamiliar name (“dax”), they should show an increased tendency to “double-check” the familiar option (as they mentally reject it: “that is a ball, so it cannot be a dax”)) before bringing their eyes back to and pointing at the unfamiliar target as the winner. In contrast, fewer double-checks are predicted on *Target-Fixated* familiar name trials (e.g., looking at the ball while hearing “ball”), where a Disjunctive Syllogism is *not* required; on these trials, children should tend to remain fixated on the target (ball) and immediately point, as they do not need to mentally exclude the other object to identify the ball. Similar predictions held for the other two tasks (Social-Learning and Explicit Negation). Figure 1 details each trial type in the three tasks (children saw multiple examples, intermixed, of each trial type with a wide variety of objects/animals).

In summary, there are at least two ways to succeed in our tasks, a Disjunctive Syllogism or a positive search for an object that contrasts with the competitor (e.g., an unfamiliar object, an unfamiliar puppet, or a non-lion animal, Figure 1). Crucially, only the Disjunctive Syllogism predicts that participants will be more likely to perform a double-check on *Target-Fixated Unfamiliar/Negation trials*. Previous research found precisely this eye-movement pattern in adults and preschoolers reasoning through the Disjunctive Syllogism (Halberda, 2006). However, it is unknown whether younger word learners deploy the Disjunctive Syllogism too, and whether this logical capacity already extends across multiple cognitive domains. In our work here with toddlers, we are predicting a developmental continuity – toddlers should be old enough to succeed on all three tasks and we expect that they (like preschoolers and adults) will succeed on *Target-Fixated Unfamiliar/Negation trials* by visually double-checking the familiar competitor in order to explicitly reject it before correctly mapping the *Unfamiliar/Negated prompt* to the target object/animal. If 2.5-year-old toddlers already have access to the domain-general Disjunctive Syllogism, then the double-checking should emerge across different tasks.

Fig. 1. Predictions of looking patterns and pointing in the three tasks. Across tasks, toddlers saw two objects/ animals that were either familiar or unfamiliar. On some trials they heard *Familiar/Assertion prompts*: a familiar name in Word Learning, a preference known to be associated with a particular animal in Social Learning, an assertion stating who was the winner in Explicit Negation. On other trials, toddlers heard *Unfamiliar/ Negation prompts*: an unfamiliar name in Word Learning, an unfamiliar preference in Social Learning, a negation stating who was NOT the winner in Explicit Negation. We predicted distinct looking patterns as a function of prompt type and whether the object fixated at the time of the prompt was the target or distractor. (A) *Target-Fixated Familiar/Assertion trials*: toddlers need not exclude the distractor to identify the target (e.g., they can recognize that "ball" refers to the ball), and thus are unlikely to shift gaze toward the distractor before pointing. (B) *Target-Fixated, Unfamiliar/Negation trials*: by hypothesis, toddlers must mentally exclude the distractor to infer the target (e.g., "dax" does not refer to the ball; therefore, "dax" must refer to the unfamiliar object). Thus, they are more likely to shift gaze to the distractor and then back to the target before pointing. (C) *Distractor-Fixated, Familiar/Assertion trials*: toddlers can directly mentally exclude the fixated distractor (e.g., "ball" does not refer to the unfamiliar object); thus, they are likely to shift gaze straight to the target, then point at it right away. (D) *Distractor-Fixated, Unfamiliar/Negation trials*: toddlers can directly mentally exclude the fixated distractor and infer the target (e.g., "dax" does not refer to the ball; therefore, it must refer to the unfamiliar object). By hypothesis, they are likely to shift gaze to the target, then point right away.

Method

Participants and Trial Exclusions. The sample size for these experiments ($N = 12$, each) was chosen based on a previous experiment with preschoolers that measured eye movements (i.e., the tendency to double-check the competitor) to investigate the Disjunctive Syllogism in word learning (Halberda, 2006). The preschooler experiment yielded a large effect size $d = 0.98$. A power analysis (cite G*Power) indicated that 12 participants would be sufficient to provide $>80\%$ statistical power to detect this large effect. Additional pilot studies in the lab suggested that toddlers would be near ceiling at pointing to the correct answer across all trial types in our tasks. Given the consistency of our methods across our three tasks, it is also possible to consider our tasks to be three conditions within a single experiment with 36 participants. Lastly, please note, this study was designed to test for a continuity in development, not to investigate developmental change.

This study was approved by the Johns Hopkins University Homewood Internal Review Board and the Harvard Internal Review Board. Parents gave written informed consent, and toddlers received a small gift (T-shirt, stuffed animal, book) for their participation.

Individual test trials were excluded from analysis for: looking at the caregiver before answering (less than 1%), fussiness (less than 1%), reaching for an object before labeling (2%), reaching for both objects (less than 1%), refusing to answer the prompt (less than 1%), experimenter's error (1.5%) or asking for a specific animal to win the race (less than 1%).

Twenty-four additional children were tested across the three experiments, but not included in the final sample due to fussiness (10), repeated refusal to point (3), equipment failure (e.g., participant gaze was not visible in the recordings (5)), or experimenter error (e.g., the experimenter used incorrect prompts or trial order, (6)). Additional participants (10) completed the task, but were excluded from the analyses since they lacked valid trials in at least one of the four conditions (*Target-Fixated, Familiar/Assertion; Target-Fixated, Unfamiliar/Negation; Distractor-Fixated, Familiar/Assertion; Distractor-Fixated, Unfamiliar/Negation*) that was a prerequisite for the double-check analysis. These happened by chance because the experimenter exerted no active control on which object a participant was fixating at the time of the prompt. Importantly, success in the pointing results remains unchanged if all of these children are included in the sample (N=46, see supplementary information).

Procedures. For all three tasks, two objects were presented on each trial, and a verbal prompt cued toddlers to choose one object (Figure 1, see Supporting Methods for the details of the three tasks' procedure). We focused on two factors. First, the target object could be indicated by either a *Familiar/Assertion prompt* or an *Unfamiliar/Negation prompt*. Second, by their own free looking, toddlers either happened to be looking at the target object or at the distractor object at the moment of hearing the prompt. This resulted in four trial types of interest, which are documented for each task in Figure 1. On approximately half of all trials, two familiar items were presented (e.g., spoon and car). When a novel item was included on a trial, approximately half of the time it was the target and half of the time it was the distractor.

First, we briefly sketch the trial structure for both *Familiar/Assertion trials* (which served as an internal control because no Disjunctive Syllogism was required) and *Unfamiliar/Negation trials* (which required Disjunctive Syllogism). For *Familiar/Assertion trials*, see row A in Figure 1. In the Word-Learning task, toddlers are already familiar with the word “ball”, and so can readily point to the ball when told “Point at the ball.” In the Social-Learning task toddlers were already familiar with panda and had previously learned that panda likes soccer (from the training trials), and so can point at the panda when asked who likes soccer. And in the Explicit Negation task toddlers simply heard the assertion that “the winner is the lion”. Importantly, none of these trial types required toddlers to reason using Disjunctive Syllogism.

For examples of *Unfamiliar/Negation trials*, see row B in Figure 1. In the Word-Learning task toddlers saw a familiar object (e.g., ball) and a totally novel object, and were asked to “Point at the [novel label, e.g., “dax”]. In the Social-Learning task toddlers saw a familiar animal whose preference they had already learned (e.g., panda), as well as a novel animal whose preference they were unaware of, and then were asked “Who likes [novel activity, e.g., “taking walks”]?”. And in the Explicit Negation task, toddlers heard the winner indicated via negation (e.g., “the winner is NOT the lion.”). Toddlers could solve these trial types by either positively mapping novel/negative information to the novel/constrasting item (i.e., Map-Noveltly-to-Noveltly) or by reasoning using Disjunctive Syllogism.

For all trial types, toddlers saw multiple trials with different animals, objects, or preferences.

To understand our behavioral predictions, consider the predicted looking for *Target-Fixated* and *Distractor-Fixated Familiar/Assertion trials* (row A & C, Figure 1). Toddlers were free to look at whichever object they liked during the period from the objects' appearance until the prompt was spoken. On most trials ($n = 78\%$), toddlers fixated each object during this time, and then they happened to be fixating either the target object (row A, Figure 1) *or* the distractor object (row C, Figure 1) at the moment of the verbal prompt (e.g., "point at the ball"). When toddlers happened to be fixating the target (e.g., ball) at the time of the prompt (row A, Figure 1), we predict that many of the trials will involve zero shifts in gaze. That is, toddlers can recognize the target and can directly point at it (in fact, it doesn't matter so much if they occasionally shift gaze on these trials, because our prediction is really a relative one – that there will be more shifts in gaze on *Target-Fixated Unfamiliar/Negation trials* (row B, Figure 1) than on *Target-Fixated Familiar/Assertion trials* (row A, Figure 1)).

On the *Distractor-Fixated Familiar/Assertion trials* (row C, Figure 1), we predict one shift in gaze as children will reject the distractor as being the target of the prompt (e.g., "point at the ball"), shift fixation to the target object (e.g., ball) recognize it and point (0→1 in row C, Figure 1).

Next, consider those trials where the target prompt is *Unfamiliar/Negation* (e.g., point at the dax), and when toddlers happen to be fixating the distractor object (e.g., ball) at the moment of hearing the prompt (row D, Figure 1). We predict one shift in gaze (i.e., away from the distractor object, towards the target object, followed by a point; 0→1 in row D, Figure 1). This is because the toddlers can reject the distractor based on the prompt, and having rejected the distractor they can infer that the other object (e.g., the novel object) must be the target (that is, they can perform a Disjunctive Syllogism).

What about those *Unfamiliar/Negation trials* where toddlers happen to be fixating the target (e.g., novel object) at the time of hearing the prompt (row B, Figure 1)? Can they directly and positively decide that the fixated object must be the target? This would be to positively map-novelty-to-novelty and to not perform a Disjunctive Syllogism. We predict instead that, on these *Target-Fixated Unfamiliar/Negation trials*, toddlers will perform a Disjunctive Syllogism and will reject the distractor object (e.g., ball) as a possible target for the prompt. We predict that the mental act of rejecting the distractor object will be linked to an increased likelihood of making a shift in gaze to it and then back (0→1→2 in row B, Figure 1). Note, a shift in gaze is not required on all of these trials, because the toddler might simply carry out the Disjunctive Syllogism in their minds, without a shift in gaze (e.g., *The winner is not the lion; he is over there; so the winner must be this bird I am looking at*). As we noted above, our prediction is a relative one: that toddlers will be more likely to make this shift in fixation on *Target-Fixated Unfamiliar/Negation trials* (row B, Figure 1) than on *Target-Fixated Familiar/Assertion trials* (row A, Figure 1). Such a difference will serve as evidence that toddlers have succeeded by working through a Disjunctive Syllogism.

Results

To probe whether preschoolers can successfully infer the target of the experimenter prompt in both the *Familiar/Assertion* and in the *Unfamiliar/Negation* condition of each task, we computed an average for each test condition (*Familiar/Assertion* vs. *Unfamiliar/Negation*), across multiple trials, and also ran Sign tests to compare the proportion of correct choices against chance (50%). We found that children performed near ceiling on all trial types across the three tasks: Word Learning, $M_{\text{Unfamiliar-Name}} = 84\%$ correct, $M_{\text{Familiar-Name}} = 96\%$; Social Learning, $M_{\text{Unfamiliar-Preference}} = 95\%$, $M_{\text{Familiar-Preference}}$

= 95%; Explicit Negation, $M_{Negation} = 95\%$, $M_{Assertion} = 98\%$. These values were all well above chance (all two-tailed Sign test p 's < .001, $n = 12$).

Toddlers' double-check analysis.

Our main question is not whether the children would succeed in the three tasks. Indeed, we have plenty of evidence that 2.5-year-old toddlers can find the referent of novel names (e.g., Halberda, 2003; Spiegel & Halberda, 2011; Lewis et al., 2020) or respond to verbal negation (Austin et al., 2014; Feiman et al., 2017). Our target question is what computations support this success across the different tasks: do toddlers use the Disjunctive Syllogism in the three tasks, as adults do in Word Learning and Explicit Negation tasks (see Halberda, 2006). To test the predictions of our hypothesis of continuity in logical development (see Figure 1), we coded whether exactly one double-check was executed on trials where toddlers chose correctly during the decision phase (i.e., the temporal window starting from the end of the experimenter's question until the child pointed to or grasped one of the two objects).

An appropriate statistical approach to answer our question is to fit mixed effects logistic regression models predicting the presence of the double-check in the *Target-Fixated trials* ($N=196$), using the *r* package *lm4*. Each model had a fixed effect for Task (i.e., Word-Learning vs. Social-Learning vs. Explicit-Negation) and a random intercept for participant, nested in Task. The first model "H0" had no additional effects. The second model "H1" had a fixed effect for Prompt Type (i.e., *Familiar/Assertion* vs. *Unfamiliar/Negation*). The third model "H2" had a fixed effect for Prompt Type and an interaction effect between Prompt Type and Task. To compare pairs of models we computed a Bayes factor (Linzen & Gallagher, 2017), which is a ratio of the likelihoods of the observed data under two different models. The Bayes Factor between two models was approximated by the Bayes Information Criteria of the models (Wagenmakers, 2007) (see supplementary code for details about the analysis). The Bayesian analysis of double-checks on *Target-Fixated trials* revealed that the data are 51 times more likely under H1 than H0: in the *Target-Fixated trials* the double-check is more probable in the *Unfamiliar/Negation* condition than in the *Familiar/Assertion* one, (*Unfamiliar/Negation* = 59%; *Familiar/Assertion* = 26%). Furthermore, the analysis revealed that the data are 143 times more likely under H1 than H2: that is, there is no interactions and the three tasks resulted in a similar pattern of results (see Fig. 2). This result was confirmed by an ANOVA (Type II Wald chi square tests, *r* package *cran*) that detected a main effect of Prompt Type ($X^2(1, N = 196) = 21.07, p < .001$), a main effect of Task (*Word-learning* = 59%; *Social-Learning* = 26%; *Explicit-Negation* = 26%; $X^2(2, N = 196) = 12.88, p < .01$), but no interaction effect ($X^2(2, N = 196) = 0.59, p > .05$). Planned follow-up pairwise comparisons (package *r emmeans*) corroborated the effects of Prompt Type in each Task: *Word-learning* (*Unfamiliar*, $M = 40\%$, $SE = 9\%$; *Familiar*, $M = 8\%$, $SE = 4\%$; $p < .01$), *Social-learning* (*Unfamiliar*, $M = 70\%$, $SE = 8\%$; *Familiar*, $M = 37\%$, $SE = 6\%$; $p < .01$), *Explicit-Negation* (*Unfamiliar*, $M = 65\%$, $SE = 8\%$; *Familiar*, $M = 30\%$, $SE = 9\%$; $p < .05$).

Next, we considered the scanning pattern that occurred on *Distractor-Fixated trials* both *Familiar/Assertion* and *Unfamiliar/Negation*. On these trials ($N = 184$), by our hypothesis, children should have no reason to double-check. They should simply reject the distractor, move their eyes to the target, and indicate their choice. Crucially, by our hypothesis there should be no difference in the probability of double-checks on *Unfamiliar/Negation trials* versus *Familiar/Assertion trials* in cases where children were fixating the distractor. To test this, we applied the same Bayesian analysis to the *Distractor-Fixated trials*. The Bayesian analysis revealed that the data are 2.4×10^3 times more likely under H0 than H1: that is, on *Distractor-Fixated trials*, there was no effect of Prompt Type on the probability of observing double-checks

(*Unfamiliar/Negation* = 2%; *Familiar/Assertion* = 2%). The data were 70 times more likely under H1 than H2. There were no interactions – the three tasks yielded a similar pattern of results (see Fig. 2).

Fig. 2. Percentage of trials with one double-check as a function of prompt type and object fixated. Bars reflect percentage of trials with one double-check prior to object choice. Across all three tasks, the percentage of double-checks was higher when toddlers were fixating the target and heard an *Unfamiliar/Negation* prompt, consistent with Disjunctive Syllogism. Bayes Factor tests reveal that an effect of prompt type on *Target-Fixated* trials is >50 times *more* likely than no effect (*), and that an effect of prompt type on *Distractor-Fixated* trials is >50 times *less* likely than no effect (†)

Toddlers' redundant saccades analysis.

If toddlers were indeed drawing a Disjunctive Syllogism on *Unfamiliar/Negation* trials, then they should (ideally) be confident in their decision. That is, if toddlers construe each trial as a 2-alternative forced-choice, having moved their attention to the distractor object and rejected it as a possible target, toddlers can move directly to the target and choose it, without further inspection of the two objects. One form of data that might index the level of toddlers' confidence post-elimination of the alternative is children's probability of making *redundant eye movements* between the objects before pointing – after having had the opportunity to check the distractor and move to the target. We predicted that, if toddlers work through a logical inference and have high-confidence in their decision, redundant saccades would *not* be higher in the *Unfamiliar/Negation* condition than in the *Familiar/Assertion* condition.

To test this prediction, we scored the presence of redundant saccades in a trial if, during the decision phase, a participant performed one or more saccades between the two potential targets above those predicted by our account (see Figure1).

We compared the probability of redundant saccades in the *Distractor-Fixated Familiar/Assertion trials* with the *Unfamiliar/Negation* ones, with the same Bayesian analysis used for the double-check. By our hypothesis there should be no difference in the probability of redundant saccades in the two conditions. Indeed, the analysis revealed that the data are 7.7×10^2 times more likely under H0 than H1: in the *Distractor-Fixated trials*, there is no effect of Prompt Type on the probability of redundant saccades (*Unfamiliar/Negation* = 8%; *Familiar/Assertion* = 6%). Furthermore, the analysis revealed that the data are 1.4×10^2 times more likely under H1 than H2: there is no interactions and the three tasks resulted in a similar pattern of results.

We also compared the probability of redundant saccades in the *Target-Fixated trials* by the same analysis. The analysis revealed that the data are 7.6×10^3 times more likely under H1 than H0: in the Target Fixated trials, redundant saccades are *more likely* in the *Familiar/Assertion* condition than in the *Unfamiliar/Negation* one (*Unfamiliar/Negation* = 2% ; *Familiar/Assertion* = 31%). Furthermore, the analysis revealed that the data are 1.5×10^2 times more likely under H1 than H2: there is no interactions and the three tasks resulted in a similar pattern of results.

Discussion

Toddlers robustly succeeded in inferring the name of a novel object, the preference of a novel character and the winner of a race based on explicit verbal negation, in line with previous findings that even infants can find the referent of a novel name (e.g., Halberda, 2003; Pomiechowska et al., 2021), and that toddlers respond to verbal negation (Austin et al., 2014; Feiman et al., 2017; Grigoroglou et al., 2019). It was however unknown whether in doing so they deployed a logical inference (Mody & Carey, 2016; see also Leahy and Carey, 2020), as adults do (see Halberda, 2006).

Toddlers' spontaneous eye-movements in the three experiments revealed a positive answer to our query. For all three tasks, on the critical *Target-Fixated Unfamiliar/ Negation trials*, children were significantly more likely to visually double-check the competitor before choosing the target object (e.g., novel object), compared to trials when they were fixating the target (e.g., ball) but hearing a *Familiar/ Assertion prompt* (where a Disjunctive Syllogism is unnecessary; Figure 2). This behavioral pattern supports the inference that toddlers performed a Disjunctive Syllogism on *Unfamiliar/ Negation trials*. Toddlers also performed few double-checks on trials where they were fixating the distractor at the time of the prompt, demonstrating that when double-checks are unnecessary (because children are already fixating the to-be-rejected distractor), children do not tend to perform them. Additional analyses showed that children also rarely made redundant saccades; revealing that children are confident in their conclusions, and that hearing *Unfamiliar/ Negation prompts* does not simply lead to a pointless increase in gaze shifts. Thus, across tasks and trials, toddlers' looking behavior was consistent with them making a Disjunctive Syllogism on precisely those trials where it would be helpful. The predictions generated by hypothesizing continuity in logical development were born out, indicating that, before robust evidence of logical vocabulary use for "or" (Chierchia et al. 2001; Morris, 2008; Singh et al. 2016; Mody and Carey 2016), 2.5-year-old toddlers used the Disjunctive Syllogism in each of our three tasks, like adults and older children do.

Our findings show that toddlers can represent two alternatives in a disjunctive relation (i.e., A OR B) and then infer the correct one by excluding the other (i.e., NOT A; THEREFORE, B). Might toddlers' responses in each of our three tasks emerge from *a single, domain-specific competence*? This is unlikely, since to answer correctly in each experiment they needed to rely on highly heterogeneous kinds of competencies: excluding that an unknown name refers to a known object, excluding that a puppet holds potentially conflicting preferences, and understanding that the word “not” was used to deny the victory of one of the competitors in a race. Might our three tasks draw instead upon *three different, task-specific, understandings of Disjunctive Syllogism*? It seems implausible that toddlers succeeded in the Explicit Negation game via a computational system dedicated to identifying the winner of a race by excluding the loser. This condition alone is likely to require general-purpose reasoning/understanding. The success we observed in this task and in the other two entirely unrelated content domains, and the strikingly similar pattern of children's eye movements across all of these, indicate that children reasoned in the same way across very different contexts.

Further research will be indispensable for uncovering the origins of the broad class of logical abilities that help get thinking off the ground. Experience, maturation and the acquisition of the logical language may play a critical role to empower adult-like learning, planning and decision making. Here we uncovered strong evidence of continuity in the development of logic: the ability to reason via at least one domain-general logical computation (i.e., Disjunctive Syllogism) appears to be in place by the third year of life.

Data Availability

The data and analysis code that support the findings of this study are available in the OSF repository: xxx.

References

- Austin, K., Theakston, A., Lieven, E., & Tomasello, M. (2014). Young children's understanding of denial. *Developmental Psychology*, 50(8), 2061–2070. <https://doi.org/10.1037/a0037179>
- Cesana-Arlotti, N., Kovács, Á. M., & Téglás, E. (2020). Infants recruit logic to learn about the social world. *Nature Communications*, 11(1), 5999. <https://doi.org/10.1038/s41467-020-19734-5>
- 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. <https://doi.org/10.1126/science.aao3539>
- Feiman, R., Mody, S., Sanborn, S., & Carey, S. (2017). What Do You Mean, No? Toddlers' Comprehension of Logical “No” and “Not.” *Language Learning and Development*, 13(4), 430–450. <https://doi.org/10.1080/15475441.2017.1317253>
- Fodor, J. A. (1979). *The Language of thought* (1st paperback printing). Harvard Univ. Press.
- Gautam, S., Suddendorf, T., & Redshaw, J. (2021). When can young children reason about an exclusive disjunction? A follow up to. *Cognition*, 207, 104507. <https://doi.org/10.1016/j.cognition.2020.104507>
- Grigoroglou, M., Chan, S., & Ganea, P. A. (2019). Toddlers' understanding and use of verbal negation in inferential reasoning search tasks. *Journal of Experimental Child Psychology*, 183, 222–241. <https://doi.org/10.1016/j.jecp.2019.02.004>

- Halberda, J. (2006). Is this a dax which I see before me? Use of the logical argument disjunctive syllogism supports word-learning in children and adults. *Cognitive Psychology*, 53(4), 310–344. <https://doi.org/10.1016/j.cogpsych.2006.04.003>
- Linzen, T., & Gallagher, G. (2017). Rapid generalization in phonotactic learning. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 8(1), 24. <https://doi.org/10.5334/labphon.44>
- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, 20(2), 121–157. [https://doi.org/10.1016/0010-0285\(88\)90017-5](https://doi.org/10.1016/0010-0285(88)90017-5)
- Mervis, C. B., Golinkoff, R. M., & Bertrand, J. (1994). Two-Year-Olds Readily Learn Multiple Labels for the Same Basic-Level Category. *Child Development*, 65(4), 1163. <https://doi.org/10.2307/1131312>
- Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, 154, 40–48. <https://doi.org/10.1016/j.cognition.2016.05.012>
- Spiegel, C., & Halberda, J. (2011). Rapid fast-mapping abilities in 2-year-olds. *Journal of Experimental Child Psychology*, 109, 132–140
- Tarski, A., & Corcoran, J. (1986). What are logical notions? *History and Philosophy of Logic*, 7(2), 143–154. <https://doi.org/10.1080/01445348608837096>
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779–804. <https://doi.org/10.3758/BF03194105>

Supporting Methods

In the *Word Learning task*, children were given four warm-up trials. The experimenter placed two familiar objects simultaneously on the table (e.g., a brush, cup), then ducked his head behind the screen and asked “Can you point at the ____ [e.g., brush]?”. The trial ended when children pointed to or grasped the requested object. If children did not respond, the experimenter raised his head and repeated the question. Next, children were presented with 12 test trials. These were structured like warm-up trials, except they could come in three types. In Familiar-Name trials children either saw two familiar objects (4 trials) or one familiar object and one unfamiliar object (4 trials), and the experimenter asked them to point to a familiar object using its familiar name. On these trials, children could use their knowledge of the name’s meaning to directly identify the target object. In Unfamiliar-Name trials (4 trials), children saw one familiar and one unfamiliar object, and the experimenter asked them to point at the unfamiliar object using an unfamiliar name. On these trials, the child did not know the referent of the unfamiliar name but could identify it by reasoning disjunctively (e.g., “dax” must refer to one object or to the other; “dax” does not refer to the ball; therefore, “dax” must refer to the unfamiliar object). No object name was ever repeated (whether as warm-up or test), and no more than two consecutive trials were of the same type. The side of correct response varied pseudo-randomly.

In the *Social Learning task*, children were introduced to three stuffed animals (e.g., Elephant, Frog, Tiger) and their distinctive preferences. The experimenter began by introducing Elephant in the first-person (e.g., “Hi, I’m Elephant, nice to meet you!”). Next, Elephant told children something that he liked-- either a food or an activity (e.g., “Do you know what I like? I like soccer!”). This process of introducing the animal and its food or activity preference was repeated for Frog. Next the experimenter brought out Frog and Elephant simultaneously, then ducked his head behind the screen and stated aloud either Elephant or Frog’s preference (e.g., “I like soccer!”). He then picked his head up from behind the screen and asked children, “Who said that?” If children responded correctly, the experimenter said, e.g., “That’s right! Elephant does like soccer!” If children responded incorrectly, the experimenter said, e.g., “No, Frog doesn’t like soccer; Elephant likes soccer!” Next, the experimenter brought out a third stuffed animal who told children what they liked. These original three *Familiar Animals* were used throughout the testing session-- at least one Familiar Animal was used on each subsequent trial. The experimenter repeated practice trials like the one described above using different combinations of the Familiar Animals, keeping his head behind the screen, until it was determined that children had successfully learned what each of them liked (as defined by getting at least four consecutive trials correct; $M = 8.4$ trials to reach criterion; range: 5 to 14 trials). After reaching criterion, each child was presented with 14 test trials. These were structured like the practice trials, except they could come in three types. In Familiar-Preference trials children either saw two Familiar Animals, or one Familiar Animal and one Unfamiliar Animal, and the experimenter stated an animal’s familiar preference (e.g., “I like soccer!”) and asked children which animal had said that. On these trials, children could use their knowledge of the preference-puppet mapping (i.e., the familiar puppet’s distinctive preference) to identify who was speaking. In Unfamiliar-Preference trials (at least three trials) children saw one Familiar and one Unfamiliar animal, the experimenter stated an animal’s unfamiliar preference (e.g., “I like taking walks!”), and children were asked which animal had said that. On these trials, children did not know the owner of the unfamiliar preference and had to identify it by reasoning disjunctively (e.g., EITHER Elephant OR Bear likes pizza; but Elephant likes soccer, NOT pizza (assuming each puppet has exactly one

distinctive preference); THEREFORE, the Bear must like pizza. No Unfamiliar Animal was presented more than once, no more than three consecutive trials were of the same type, and the side of correct response varied pseudo-randomly.

In the *Explicit Negation task*, children were introduced to two stuffed animals (e.g., Frog, Tiger), and the experimenter explained that they were going to race to reach a finish line (a line of colored tape, near the child, on the table), but that children needed to help by making the winning animal cross the finish line. The experimenter then ducked his head behind the screen, said, “ready, GO!” and started to move the animals toward the finish line (i.e., towards the child). Before they crossed it, paused and said, “And the winner is the ____ [e.g., Tiger]”, and then said, “Can you help me make the winner finish the race?” (this explicit request was made only on the first trial, which was considered a practice trial). The trial ended when children pointed at or grasped one of the animals. Practice trials were repeated using different combinations of stuffed animals until it was determined that children understood the game ($M = 2.2$ trials; range: 1-4 trials). Test trials were structured like practice trials, except they could come in two types. On Assertion trials the experimenter indicated the winner by saying “And the winner is the ____ [e.g., Elephant]!” On these trials, children could directly identify the winner because its identity was explicitly stated by the experimenter. On Negation trials the experimenter indicated the winner by saying “And the winner is NOT the ____ [e.g., Lion]!” On these trials, children could identify the winner reasoning disjunctively (i.e., EITHER Elephant OR Lion will win the race. Lion is NOT the winner; THEREFORE, Elephant must be the winner). Children were presented with up to 12 test trials, depending on their interest in the task as judged by the experimenter ($M = 9.25$ trials; range: 7-12 trials). Half were Negation trials and half Assertion trials. No more than two consecutive trials were of the same type (Assertion vs. Negation) and the side of the correct response varied pseudo-randomly.