Title
Young children adapt their search behavior for necessary versus merely possible outcomes

Permalink
https://escholarship.org/uc/item/2c23r0b4

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors
Andreuccioli, Luisa
Mazor, Sophie
Begus, Katarina
et al.

Publication Date
2024

Young children adapt their search behavior for necessary versus merely possible outcomes

Luisa Andreuccioli (landreuccioli@ucsd.edu)
Department of Psychology, University of California San Diego

Sophie Mazor (smazor@ucsd.edu)
Department of Psychology, University of California San Diego

Katarina Begus (khb@psy.ku.dk)
Department of Psychology, University of Copenhagen

Elizabeth Bonawitz (elizabeth_bonawitz@gse.harvard.edu)
Graduate School of Education, Harvard University

Stephanie Denison (stephanie.denison@uwaterloo.ca)
Department of Psychology, University of Waterloo

Caren M. Walker (carenwalker@ucsd.edu)
Department of Psychology, University of California San Diego

Abstract

Although even infants appear to consider multiple possibilities, preschoolers often fail tasks that require reasoning about mutually exclusive alternatives. We review two explanations for this failure: (1) children have a minimal representation of possibility and fail to distinguish necessary from merely possible outcomes; and (2) children are sensitive to this distinction, but competing motivations (e.g., the tendency to explore) can lead to apparent failures. To test these hypotheses, we assessed 3- and 4-year-olds on a novel search task. Here, children searched for an object that was dropped from either a transparent (one necessary location) or opaque (two possible locations) set of inverted Y-shaped tubes. In Exp. 1, we found that children spent less time searching the first location when there were two possible candidates. Exp. 2 replicates these results in a digital task that does not require manual search.

Keywords: cognitive development; modal reasoning; knowledge representations; alternative possibilities

Introduction

The ability to reason about multiple possibilities is central to human cognition. Not only is this key to everyday decision-making (Fodor, 1975; Simon, 1955), it is also fundamental to the process of learning: we consider possibilities when evaluating explanations (Lombrozo, 2006), inferring causal relationships (Gopnik & Wellman, 2012), thinking counterfactually (Hegarty, 2004; Kahneman & Tversky, 1981) or planning for the future (Beck et al. 2006), and when making logical inferences from initial premises (Johnson-Laird et al., 1992; Johnson-Laird, 2010). However, there is disagreement about when and how this ability emerges in childhood.

On some accounts, the ability to represent alternative possibilities is proposed to be relatively late-emerging. According to these accounts, this capacity is strictly dependent on the acquisition of modal logic – symbolic operators that allow the learner to distinguish what is *necessary* from what is merely *possible* (Leahy & Carey, 2020). Language studies suggest that children only start comprehending such concepts around their fourth year (e.g. 'may' vs. 'have to' Ozturk & Papafragou, 2015; Leahy & Zalnieriunas, 2021).

In line with this proposal, research shows that children under four years of age appear to consistently fail tasks that require them to explicitly reason about alternative possibilities (Mody & Carey, 2016; Leahy, 2023; Leahy et al. 2022; Rohwer, Kloo, & Perner, 2012; Kim et al., 2016; Redshaw & Suddendorf, 2016). For example, in a task where children are asked to pick between a singleton cup that is guaranteed to contain a reward, and two cups that are merely possible locations, 3-year-olds only picked the 'safe bet' about half the time. This has been interpreted as evidence that young learners are *unable* to distinguish between necessity and possibility (Mody & Carey, 2016; Leahy & Carey, 2020). The same behavior is also observed in modified versions of the task that maintain the same logic: one in which working memory demands were reduced (Leahy, 2023), and one involving additional scaffolding (Leahy et al. 2022, but see Alderete & Xu, 2023 for a simplified version where children do succeed). Indeed, 3-year-olds have shown related failures in a variety of tasks that require that

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*

they prepare for mutually exclusive outcomes (e.g., Beck et al., 2006; Leahy, 2024; Redshaw & Suddendorf, 2016; Robinson et al., 2006; but see Turan-Kücük & Kibbe, 2024 for an alternative explanation of these results).

In light of this evidence, Leahy and Carey (2020) propose that, prior to age 4, children are unable to deploy full possibility concepts. They argue that children's failures (and successes) can be entirely explained by a *minimal*, as opposed to *modal* representation of possibility. While a modal learner can simultaneously hold multiple possibilities in mind, a minimal representer only considers one 'possibility' at a time. According to this account, young children and non-human animals lack a symbolic marker that allows them to conceptualize an event as merely possible. Therefore, when faced with multiple possibilities, minimal representers simulate just one outcome and treat this simulation as if it were a matter of fact (Leahy & Carey, 2020), until and unless such simulation proves to be false.

In contrast, a growing body of evidence suggests the capacity to engage in modal reasoning may be much earlier developing (Alderete & Xu, 2023; Cesana-Arlotti et al. 2018; Cesana-Arlotti et al. 2022; Goupil et al. 2016; Téglás & Bonatti, 2016; Téglás et al., 2007; Turan-Küçük & Kibbe, 2024; Xu & Garcia, 2008). For instance, 12-month-olds succeed at a non-verbal task that requires them to represent two mutually exclusive possibilities in the initial premise of a deductive inference (Cesana-Arlotti et al. 2018). Other evidence shows that infants' pupils dilate when the number of possible outcomes increases (Cesana-Arlotti et al. 2022), suggesting that they are (at least implicitly) representing possibilities. Furthermore, in a study where 18-30 month-olds observed ambiguous evidence consistent with two, equally likely hypotheses (i.e., either the blue cube activates a toy or 'different' pairs activate a toy; Goddu et al., 2021), toddlers flexibly selected between these possibilities, apparently generating and holding both options in mind.

According to Leahy and Carey (2020), much of this evidence for early competence can be construed as resulting from a minimal representation. They offer a plausible account by which children could succeed in these contexts by generating sequential simulations of a single outcome at a time (Leahy & Carey, 2020). But before accepting the minimal account, it is necessary to rule out alternative explanations for children's prior failures (see also Turan-Kücük & Kibbe, 2024).

Here, we consider the proposal that young children's performance (or failure) in tasks that require them to select a necessary outcome over merely possible ones may be driven, at least in part, by a competing motivation to explore (to seek additional information) about uncertain events. Indeed, a large body of work shows that young learners tend to explore when their knowledge is uncertain, insufficient, or incompatible with their observations (Liquin & Lombrozo, 2020; Schulz & Bonawitz, 2007; Perez & Feigenson, 2022; Lapidow et al., 2021, Wang et al, 2021). In fact, there is evidence to suggest that children will forgo the opportunity to obtain a tangible reward in order to gain new information about the world (Lapidow & Walker, 2020; Liquin & Gopnik, 2022). For example, in a classic 'multi-armed bandit' task where participants could either exploit and maximize immediate rewards, or explore at a cost, children as young as 3 were more likely to explore than older children and adults (Liquin & Gopnik, 2022). It is therefore plausible that tasks that present children with a choice between a necessary and merely possible outcome might similarly promote (costly) exploration behavior. If so, some of children's prior failures may reflect their greater motivation to learn about (rather than avoid) a merely possible outcome.

In the current studies, we remove the requirement that children pick *either* the necessary *or* a possible location, avoiding the potential for conflicting motivations to influence children's performance. Instead, we designed a novel task in which we replace this forced choice with a continuous measure of children's search behavior to assess whether even young learners' search strategy differs when responding to distinct modal concepts.

Experiment 1

In Exp. 1, we presented 3- and 4-year-old children with a fully opaque and fully transparent set of inverted Y-shaped tubes, each emptying into two opaque search boxes (Fig. 1). In a within-subject design, all children observed two trials. In each trial, a target object was dropped into one of the two sets of tubes and children were asked to locate the object inside the search boxes.

When the target object was dropped into the transparent tubes, children could observe where it landed and direct their search to the correct box to retrieve the object. When it was dropped into the opaque tubes on the other hand, children had no information about where the object fell, leading to a situation in which there was more than one possible location. Importantly, unbeknownst to the participant, all search boxes were designed to trap the target object before it landed inside. This allowed us to compare the amount of time children were willing to spend searching in each location in each set of tubes.

If children hold only a minimal representation of possibility, they should simulate a single outcome in the opaque trial, and treat that simulation as knowledge. We should therefore find no significant difference in their search behavior between the opaque and transparent trials, since the location of the object is "known" in both cases. If children are instead able to represent both search boxes as merely possible locations in the opaque trial, they should spend significantly less time searching the first location they approach in this condition compared to when the object is dropped from the transparent tube.

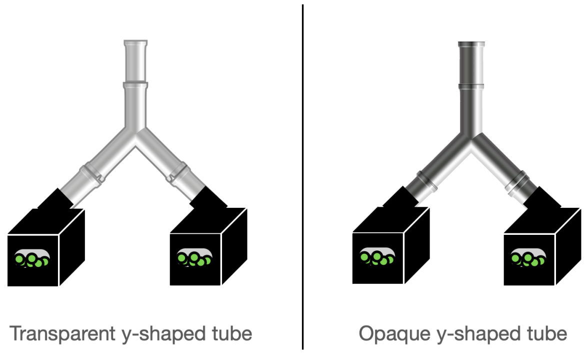

Figure 1: Illustration of the two sets of Y-shaped tubes (with search boxes attached), presented as within-subject trials in Experiment 1.

Participants

A total of 54 children were tested, including 28 3-year-olds (16 females, mean age = 3.5) and 26 4-year-olds (10 females, mean age = 4.5). This sample size was roughly based on the sample used by Leahy and colleagues (2022). Children were either monolingual English speakers or bilingual speakers with English as a primary language. An additional 10 children (9 3-year-olds and 1 4-year-old) were tested, but excluded (1 withdrawal, 1 failure to attend, 1 failure on familiarization, 3 experimenter errors, 3 stimuli failures).

Materials

Materials included a 14mm yellow-painted wooden bead (the target object), stickers, and two sets of Y-shaped tubes built from PVC pipes (see Fig. 1) attached to an inclined plane. One set of tubes was opaque and the other set was transparent. The top opening was divided to allow the experimenter to control the side that the target object would fall when dropped. This divider was not visible to participants.

Six black, rectangular, cardboard search boxes were built for the study. The side of the box facing the child included an opening covered by a felt flap. This allowed the child to reach in with one hand to explore the contents, without visual access. On the opposite side of the box (facing the experimenter), a hidden drawer-like opening allowed the experimenter to secretly push the target object inside the search box. The search boxes included a fabric tube extending from the top, so they could be attached to the bottom of each of the openings of the Y-shaped tubes during test trials. A layer of fabric inside the extension prevented the target object from actually landing in the boxes. Each box was filled with ten, multicolored pom poms, which ensured that children had to manually search inside of the box to find the target object.

Finally, an iPad was used as a 31-second timer. The iPad played a musical track that increased in tempo to signal that time was running out.

Procedure

The child sat across the table from the experimenter. A mitten was placed on the child's non-dominant hand to ensure they could only search inside one box at a time. Children were instructed to only use their free hand when prompted to search, and were reminded of this as needed.

Introductory phase The introductory phase was intended to familiarize children with the items involved in the manual search task. The experimenter presented one of the search boxes and told the child that it was filled with soft, fluffy balls. Children had the opportunity to touch the fluffy balls when outside and inside the box. Next, they were asked to hold the yellow wooden ball (target object), and asked to report whether it felt "soft or hard" to draw their attention to the difference between the box filler and the target. The child was told that their goal was to "find as many hard balls as possible", and that they would win a sticker every time they found one. The experimenter then dropped the target object from an opening at the top of the box and asked the child to use their dominant hand to find it and remove it.

Familiarization to search boxes The experimenter presented two new search boxes and placed them side-by-side in front of the child. The experimenter said, "There is already a hard ball inside one of these boxes, and I am going to ask you to find it. You can search this box, this box, or both boxes (touching each box in turn)". On the first familiarization trial, the experimenter covertly pushed the target object inside whichever box the child approached first.

Following this, the experimenter brought out two new search boxes and the same procedure was repeated. This time, however, the experimenter pushed the target object inside whichever box the child did *not* approach first so they would have to spontaneously switch to this alternative location to find it. If the child stopped searching before finding the target object, they were given a series of increasingly suggestive prompts. Prompts were only given as needed, and in the following order: (1) "Have you looked all around the box?" (2) "Remember, you will win a sticker if you can find the hard ball!" (3) "There is definitely a hard ball inside one of these boxes." (4) "You can use all the time you need to find it!" If children expressed frustration for not being able to find it in a box, and after using all other prompts, we moved onto (5) "Where else could it be?". This final prompt was only needed for a small minority of participants.

On the next two familiarization trials, the experimenter presented two new search boxes and repeated the same procedure. On these trials, a timer was introduced, and the experimenter instructed the child that they had to find the hard ball before the music stopped to win a sticker. Again, the target object was covertly hidden in the first location the child approached. Then, on the subsequent timed trial, it was hidden in the location that the child did not initially search. After each familiarization trial, the child was given positive or neutral feedback, depending on whether they found the ball before the timer ran out. Children had to find the target object on at least one of the four familiarization trials to be included in the final sample. Only one child failed to do so.

Familiarization to Y-shaped tubes The experimenter next introduced the transparent set of Y-shaped tubes. They explained that they would drop the ball from the top of the tubes, emphasizing that children would be able to see the ball falling down one of the tubes. The child was asked to point to the tube that the object fell from. The object was dropped four times in a pseudo-random order [L-R-R-L]. Next, the experimenter introduced the opaque set of Y-shaped tubes, drawing attention to the fact that this time it would *not* be possible to see the ball as it fell down the tube. Again, the ball was dropped four times in a pseudo-random order [R-L-L-R], and children were asked to point to the tube that produced the ball.

Test phase At test, the experimenter presented either the transparent or opaque set of Y-shaped tubes (counterbalanced) and said: "Remember those boxes we played with earlier? I am going to attach them to the ends of the tubes. These boxes only have soft balls inside them! Now, I will drop a hard ball from up here, and I am going to ask you find it."

As in the familiarization trials, the experimenter reminded the child that they had to find the hard ball before the timer ran out if they wanted to win a sticker. Critically, since the target object never actually fell into either box, the child never found the ball on either test trial. Use of the timer ensured a plausible explanation for their inability to locate the ball. Children were given neutral feedback upon not finding the ball and encouraged to try again. On the second test trial, the same procedure was repeated with the second set of tubes. For both test trials, we coded time (in seconds) spent searching the first location, time spent in each location after the child switched, and the total number of times they switched during a given trial. On the transparent trial, we also coded whether children started searching in the correct location. The two 4-year-old children who started searching in the incorrect location were still included in our analyses.

Children's search times were recorded using Datavyu (Datavyu.org), a software that allows for frame-by-frame behavioral coding. Reliability coding was completed for all trials with 99.99% agreement.

Results and Discussion

Planned statistical analyses were conducted in R (version 4.1.1). For each of the two age groups, we asked whether there was a significant difference in the time spent searching in the first location they approached in the opaque vs. transparent test trial. This revealed that 3-year-olds searched significantly longer on the transparent (M = 24.07, SD = 9.83) compared to the opaque trial (M = 13.95, SD = 11.25), $t(27)$ = 4.5638, p < 0.01. The same was true for 4-year-olds, (transparent: M = 20.40, SD = 11.22; opaque: M = 6.76 , SD = 6.19), $t(25)$ = 6.13, p < 0.01.

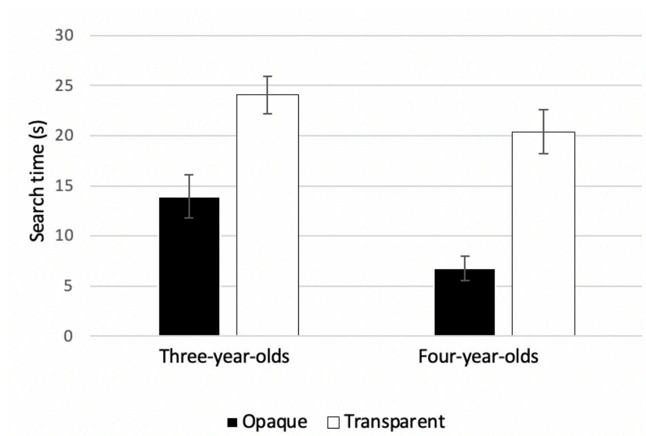

Figure 2: Three- and four-year-olds' search time in the first location approached on transparent and opaque trials in Experiment 1. Error bars represent SEM.

Three-year-olds also switched significantly more times in the opaque trial, when there were two possible locations (M = 1.29, SD = 1.03) compared to the transparent trial, where there was only one (M = 0.62, SD = 0.88), $t(26)$ = -3.1224, p < .005. The same was true for 4-year-olds (opaque: M = 2.46, SD = 1.20; transparent: M = 1.07 , SD = 1.26), $t(25)$ = -4.7108, p < .005. There was no significant difference between 3- and 4-year-olds' performance**,** p = .268.

Consistent with our hypothesis, we found that both 3- and 4-year-olds searched longer in the first location they approached when there was only one possible location (transparent tubes), compared to the case when there were two possible locations (opaque tubes). Both 3- and 4-year-olds also switched significantly more often in the opaque trial compared to the transparent. These findings provide initial evidence that, as a group, both 3- and 4-year-old children modify their search strategy, depending on whether the outcome is necessary or merely possible.

Experiment 2

In Experiment 2, we replicate these findings in a digital paradigm that did not rely on manual search. We designed an iPad version of the task that replaced manual search with lever taps. This change reduced the possibility that children's search might be guided, directly or indirectly, by sensory feedback. Although children produced discreet taps, we again recorded time spent searching (in seconds) as the dependent variable.

Participants

Participants were 24 3-year-olds (11 females, mean age = 3.5) and 24 4-year-olds (13 females, mean age = 4.5). An additional 4 3-year-olds and 2 4-year-olds were tested but excluded (3 withdrew, 3 technical issues).

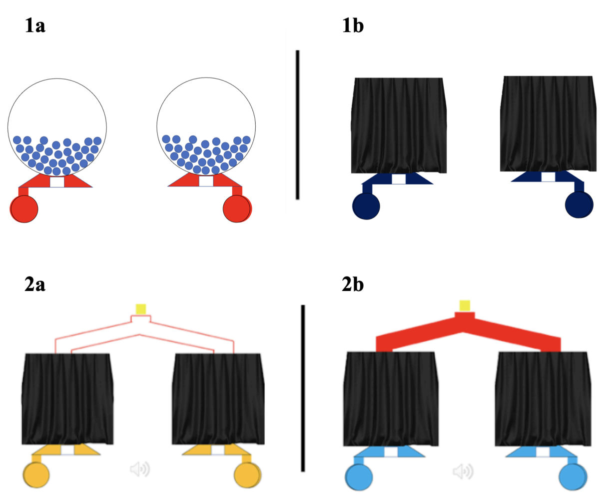

Figure 3: (1a) Introductory phase presenting two machines; (1b) Familiarization trial with two occluded machines; (2a) Test trial with transparent Y-shaped tubes; (2b) Test trial with opaque Y-shaped tubes

Materials and Procedure

Materials included a 12.9 inch display iPad Pro with a kickstand case, a Bluetooth remote control clicker, a child-size mitten, and stickers.

The study was created in PowerPoint. Search boxes were replaced with gumball machines that could be placed behind an occluder and attached to the bottom of fully transparent or fully opaque Y-shaped tubes. Each gumball machine was filled with blue marbles and had a lever that could be tapped to dispense the objects inside. Each tap corresponded to a single object being released from the opening on the bottom. Children were instructed to keep searching (tapping the lever) until the target object (a yellow square) came out.

Despite these minor changes to the framing, the general procedure remained largely the same as that used in Experiment 1, except that we reduced the number of familiarization trials from 4 to 2.

Introductory phase Children were instructed that their job would be to find as many yellow squares as they could to win stickers. The screen then showed two gumball machines filled with blue marbles (see Fig. 3). The experimenter demonstrated how the machines worked, and children were given a chance to press each machine's lever once.

Next, a new machine appeared on screen, and children watched as the target object was dropped inside. They were then asked to tap the lever until the yellow square came out. This was repeated with a new, occluded machine. Next, two new occluded machines appeared on the screen, and children were told that a yellow square had already been hidden inside one of the machines. During this trial, the square was hidden inside the left machine and was dispensed on the sixth press. When the child succeeded, they were rewarded with a sticker. Two new occluded machines appeared on the screen, and children were instructed to find the yellow square before the timer ended. This time, the target object was hidden inside the right machine and was dispensed on the third press. Just like in Exp. 1, if children found the target object before the timer ended, they were awarded another sticker.

The coding scheme was identical to the one used for Study 1. Three 3-year-olds and one 4-year-old started searching in the incorrect location on the transparent trial, and these children were still included in the analyses. We used Datavyu to record search times, and reliability coding on 25% of trials resulted in 99.99% agreement.

Results and Discussion

Again, planned statistical analyses were conducted in R (version 4.1.1). As in Experiment 1, both 3- and 4-year-olds spent significantly longer searching the first location they approached in the transparent (3s: M = 23.30, SD = 11.86; 4s: M = 26.33, SD = 9.54) compared to the opaque condition (3s: M = 15.47, SD = 12.17; 4s: M = 12.74, SD = 11.72), $t(23) = 3.1613$, $p < .002$ and $t(23) = 5.0363$, $p < .005$, respectively (see Fig. 4).

Both 3- and 4-year-olds also switched significantly more times in the opaque, compared to the transparent trial, (3s: $t(23) = -2.9132$, $p = 0.007$; 4s: $t(23) = -2.4862$, $p = .02$, and there was no significant difference between the two age groups ($p = .1$) Overall, these results replicate the findings from Exp. 1 in an iPad version of the task.

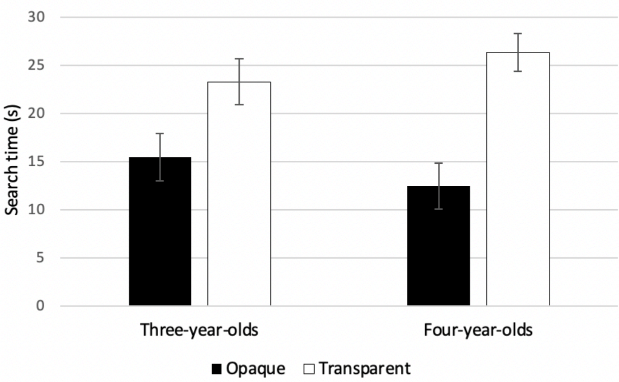

Figure 4: Three- and four-year-olds' search time in the first location approached on transparent and opaque trials in Experiment 2. Error bars represent SEM.

General Discussion

In the current studies, we investigate whether 3- and 4-year-old children can represent multiple, mutually exclusive possibilities using a novel search task. In Experiment 1, we found that children switched to the second available location significantly faster in the opaque condition (when both locations were possible), compared to the transparent condition (where only one location was possible). Children in both age groups also switched more frequently when the object was dropped from the opaque, compared to the transparent set of tubes, and there was no significant difference between 3- and 4-year-olds' performance. These results demonstrate that both 3- and 4-year-old children adapt their search strategy for necessary versus merely possible options, suggesting that they are sensitive to the presence of multiple possibilities in the opaque case. In Experiment 2, we replicate these findings using a digital version of the task that does not rely on manual search.

These findings provide initial evidence that children under four can distinguish what is necessary from what is merely possible. According to the minimal account, when confronted with multiple possible outcomes, children simulate a single outcome and treat that simulation as knowledge, until and unless their simulation is proven incorrect (Leahy & Carey, 2020). If this were the case, children in the current studies would have represented the outcome in both the transparent and opaque trials as necessary, and should not display any differences in search behavior between the two conditions. In contrast, not only did children search the first location they approached for significantly less time in the opaque case, they also switched between locations more frequently, further supporting the interpretation that they are considering multiple locations.

One interesting question to be addressed in future work is whether the differences we observe at the group level remain when examining patterns of performance in individual children. It is possible that we will find variation, particularly in 3-year-olds' ability to differentiate these concepts (see Supporting Information in Leahy et al. 2022 for a discussion of this possibility). Future work will also examine performance in younger children to better understand the developmental trajectory of this early sensitivity to modal concepts.

Although additional research is needed to determine whether children's conflicting motivation to explore may account for some of their previous failures, the current findings are compatible with this proposal. Unlike prior work, which has operationalized modal reasoning in terms of children's ability to choose a 'safe bet' over two merely possible options (e.g., Mody & Carey, 2016; Leahy et al. 2022; Leahy, 2023), children in the current studies were not forced to pick. In removing the requirement that they forgo the opportunity to explore in favor of a guaranteed reward, we find evidence that young children do differentiate between a necessary and a merely possible outcome.

Finally, an important limitation of these experiments, currently being addressed in ongoing work, concerns our operationalization of 'necessity' and 'possibility'. Specifically, since the set of transparent tubes allowed children *visual access* to the trajectory of the object, while the opaque set of tubes did not, children's representations of these two concepts were based on qualitatively different types of evidence. While their representation of necessity was based on *direct observation,* their representation of possibility was based on *inference*. Indeed, some prior work suggests that visually tracking an object may lead to a stronger representation (Call & Carpenter, 2001), which could have resulted in the relatively longer search time we observe in the transparent trial. To examine this possibility, we are now collecting data on a third experiment that uses a modified version of Experiment 2. Here, *both* representations are reached via inference: the fully transparent set of tubes is replaced with a set in which one arm is opaque, and one is transparent. In this task, the target object will always fall through the opaque (and not the transparent) branch, and children will therefore have to *infer* (rather than directly observe) that the object must have fallen through the opaque side. This ensures that the strength of each representation is matched between conditions, allowing us to rule out this potential alternative explanation.

In sum, these initial experiments provide support for the claim that sensitivity to multiple, mutually exclusive possibilities is early developing. This work also serves to introduce a novel task that can be used to assess aspects of modal reasoning in future work.

Acknowledgements

We thank the Fleet Science Center for supporting this research. We thank Brian Leahy and Susan Carey for helpful discussions. EB was funded by the McDonnell Foundation. CW was funded by the Jacobs Foundation.

References

Alderete, S., & Xu, F. (2023). Three-year-old children's reasoning about possibilities. *Cognition*, *237*, 105472. doi: 10.1016/j.cognition.2023.105472

Beck, S. R., Robinson, E. J., Carroll, D. J., & Apperly, I. A. (2006). Children's thinking about counterfactuals and hypotheticals as possibilities. Child Development, 77(2), 413–423. doi: 10.1111/j.1467-8624.2006.00879.x

Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? Animal Cognition, 3, 207-220. doi: 10.1007/s100710100078

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. doi: 10.1126/science.aao3539

Cesana-Arlotti, N., Varga, B., & Téglás, E. (2022). The pupillometry of the possible: an investigation of infants' representation of alternative possibilities. *Philosophical Transactions of the Royal Society B*, *377*(1866), 20210343. doi: 10.1098/rstb.2021.0343

Fodor. J.A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press. doi: 10.2307/2184356

Goddu, M. K., Sullivan, J. N., & Walker, C. M. (2021). Toddlers learn and flexibly apply multiple possibilities. Child Development, 92(6), 2244-2251. doi: 10.1111/cdev.13668

Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: causal models, Bayesian learning mechanisms, and the theory theory. Psychological Bulletin, 138(6), 1085. doi: 10.1037/a0028044

Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. Proceedings of the National Academy of Sciences of the United States of America, 113(13), 3492–3496. doi: 10.1073/pnas.1515129113

Hegarty, M. (2004). Mechanical reasoning by mental simulation. Trends in Cognitive Sciences, 8(6), 280-285. doi: 10.1016/j.tics.2004.04.001

Johnson-Laird, P. N., Byrne, R. M., & Schaeken, W. (1992). Propositional reasoning by model. Psychological Review, 99(3), 418. doi: 10.1037/0033-295X.99.3.418

Johnson-Laird, P. N. (2010). Mental models and human reasoning. Proceedings of the National Academy of Sciences, 107(43), 18243-18250. doi: 10.1073/pnas.1012933107

Kahneman, D., & Tversky, A. (1981). *The simulation heuristic*. Stanford Univ CA Dept of Psychology. doi: 10.1017/CBO9780511809477.015

Kim, S., Paulus, M., Sodian, B., & Proust, J. (2016). Young children's sensitivity to their own ignorance in informing others. PLoS ONE, 11(3). doi: 10.1371/journal.pone.0152595

Lapidow, E., & Walker, C. M. (2020). Informative experimentation in intuitive science: Children select and learn from their own causal interventions. Cognition, 201, Article 104315. doi: 10.1016/j.cognition.2020.104315

Lapidow, E., Killeen, I., & Walker, C. M. (2021). Learning to recognize uncertainty vs. recognizing uncertainty to learn: Confidence judgments and exploration decisions in preschoolers. Developmental science, 25(2), e13178. doi: 10.1111/desc.13178

Leahy, B. P., & Carey, S. E. (2020). The acquisition of modal concepts. Trends in Cognitive Sciences, 24(1),65-78. doi: 10.1016/j.tics.2019.11.004

Leahy, B., Huemer, M., Steele, M., Alderete, S., & Carey, S. (2022). Minimal representations of possibility at age 3. *Proceedings of the National Academy of Sciences*, *119*(52), e2207499119. doi: 10.1073/pnas.2207499119

Leahy, B. (2023). Don't you see the possibilities? Young preschoolers may lack possibility concepts. *Developmental Science*, e13400. doi: 10.1111/desc.13400

Leahy, B. (2024). Many preschoolers do not distinguish the possible from the impossible in a marble-catching task. *Journal of Experimental Child Psychology*, *238*, 105794. doi: 10.1016/j.jecp.2023.105794

Leahy, B., & Zalnieriunas, E. (2021). Might and might not: Children's conceptual development and the acquisition of modal verbs. In *Semantics and Linguistic Theory* (pp. 426-445). doi: 10.3765/salt.v31i0.5082

Liquin, E. G., & Lombrozo, T. (2020). Explanation-seeking curiosity in childhood. Current Opinion in Behavioral Sciences, 35, 14–20. doi: 10.1016/j.cobeha.2020.05.012

Liquin, E. G., & Gopnik, A. (2022). Children are more exploratory and learn more than adults in an approach-avoid task. Cognition, 218, 104940. doi: 10.1016/j.cognition.2021.104940

Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, *10*(10), 464-470. doi: 10.1016/j.tics.2006.08.004

Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. Cognition, 154, 40-48. doi: 10.1016/j.cognition.2016.05.012

Ozturk, O., & Papafragou, A. (2015). The acquisition of epistemic modality: From semantic meaning to pragmatic interpretation. *Language Learning and Development*, *11*(3), 191-214. doi: 10.1080/15475441.2014.905169

Perez, J., & Feigenson, L. (2022). Violations of expectation trigger infants to search for explanations.

Cognition, 218, 104942. doi: 10.1016/j.cognition.2021.104942

Redshaw, J., & Suddendorf, T. (2016).Children's and apes' preparatory responses to two mutually exclusive possibilities. *Current Biology*, *26*(13), 1758-1762. doi: 10.1016/j.cub.2016.04.062

Robinson, E. J., Rowley, M. G., Beck, S. R., Carroll, D. J., & Apperly, I. A. (2006). Children's sensitivity to their own relative ignorance: Handling of possibilities under epistemic and physical uncertainty. Child Development, 77(6), 1642-1655. doi: 10.1111/j.1467-8624.2006.00964.x.

Rohwer, M., Kloo, D., & Perner, J. (2012). Escape from metaignorance: How children develop an understanding of their own lack of knowledge. Child Development, 83(6), 1869–1883. doi: 10.1111/j.1467-8624.2012.01830.x

Simon, H. A. (1955). A behavioral model of rational choice. Quarterly Journal of Economics, 69(1), 99-118. doi: 10.2307/1884852

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. Developmental Psychology, 43(4), 1045–1050. doi: 10.1037/0012-1649.43.4.1045

Téglás, E., Girotto, V., Gonzalez, M., & Bonatti, L. L. (2007). Intuitions of probabilities shape expectations about the future at 12 months and beyond. Proceedings of the National Academy of Sciences, 104(48). doi: 10.1073/pnas.0700271104

Téglás, E., & Bonatti, L. L. (2016). Infants anticipate probabilistic but not deterministic outcomes. Cognition, 157, 227–236. doi: 10.1016/j.cognition.2016.09.003

Turan-Küçük, E. N., & Kibbe, M. M. (2024). Three-year-olds' ability to plan for mutually exclusive future possibilities is limited primarily by their representations of possible plans, not possible events. Cognition, 244, 105712. doi: 10.1016/j.cognition.2023.105712

Wang, J., Yang, Y., Macias, C., & Bonawitz, E. (2021). Children with more uncertainty in their intuitive theories seek domain-relevant information. *Psychological Science, 32*(7), 1147-1156. doi: 10.1177/0956797621994230

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, *105*(13), 5012-5015. doi: 10.1073/pnas.0704450105