

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Sensitivity to Ostension is Not Sufficient for Pedagogical Reasoning by Toddlers

Permalink

<https://escholarship.org/uc/item/6jm3d6hr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Tecwyn, Emma C.

Seed, Amanda M.

Buchsbaum, Daphna

Publication Date

2020

Peer reviewed

Sensitivity to Ostension is Not Sufficient for Pedagogical Reasoning by Toddlers

Emma C. Tecwyn (emma.tecwyn@bcu.ac.uk)

Department of Psychology, Birmingham City University, Curzon Building, 4 Cardigan Street, Birmingham, B4 7BD, UK

Amanda M. Seed (ams18@st-andrews.ac.uk)

School of Psychology and Neuroscience, University of St Andrews, St Mary's Quad, South St, St Andrews KY16 9JP, UK

Daphna Buchsbaum (buchsbaum@psych.utoronto.ca)

Department of Psychology, University of Toronto, 100 St George Street, Toronto, ON M5S 3G3, CANADA

Abstract

To investigate the role of ostensive cues in pedagogical reasoning, we explored whether toddlers, like preschoolers, would copy causally implausible actions following a pedagogical demonstration. Toddlers watched a demonstrator perform a two-action sequence (AB) on a puzzle-box that led to a reward. We manipulated the demonstrator's intentionality and the causal plausibility of action A and examined how these factors influenced copying behavior. Although toddlers were more likely to copy A when it was causally plausible, they were not influenced by the demonstrator's intentionality. Importantly, toddlers were no more likely to copy the AB sequence following a pedagogical demonstration vs. a non-communicative demonstration. Comparing behavioural data to computational model predictions for learners differing in their sensitivity to intentionality and causal plausibility supported an absence of pedagogical reasoning. These results suggest that sensitivity to ostension may be a necessary prerequisite—but is not sufficient for—pedagogical reasoning in a causal imitation task.

Keywords: causal reasoning; cognitive development; ostension; overimitation; pedagogy; social learning

Introduction

The ability to learn effectively from observing the actions of others is fundamental for the transmission and maintenance of uniquely-human culture (e.g., Tomasello, 1999; Boyd, Richerson & Henrich, 2011). When observing how an individual interacts with a causally-opaque object, being able to understand the intentions behind their actions can provide important information regarding which actions are necessary, and can influence the inferences a learner makes about how that object functions (e.g., Bonawitz et al., 2011; Buchsbaum et al., 2011; Shafto, Goodman & Griffiths, 2014). For example, imagine you watch someone using a machine that dispenses candy. They look at you and say “Hey, watch this!” then they first tap the top of the machine, before turning a dial, and a candy pops out. When it is your turn, what do you do? A 4-year-old in this situation would typically copy both of the actions, even though tapping the top of the machine seems causally unrelated to making the candy come out (e.g., Horner & Whiten, 2005; Lyons, Young & Keil, 2007).

Copying causally irrelevant actions—termed overimitation (Lyons et al., 2007)—can be seen as rational if a learner is capable of pedagogical reasoning—that is, if they recognize

that a demonstrator's actions are being done communicatively for their benefit (e.g., Butler & Markman, 2012). Specifically, if a demonstrator is communicative, knowledgeable and informative, Bayesian computational models have shown that this licenses a strong inference that all of the demonstrated actions should be copied—after all, why would a helpful teacher have included a particular action if it was not necessary (e.g., Bonawitz et al., 2011; Buchsbaum et al., 2011; Shafto et al., 2014)?

A recent study combining behavioral data with Bayesian computational modelling suggested that 3- to 5-year-old children distinguish pedagogical, intentional and unknowing demonstrations when deciding which actions to copy in an imitation task (Buchsbaum et al., under review). After watching a pedagogical demonstrator perform a sequence of two actions (AB) on a puzzle-box that led to a reward, 3- to 5-year-olds faithfully copied both actions, even if the first action (A) was causally implausible. In contrast, if the demonstrator was intentional but non-communicative, children were less likely to copy the implausible action. This performance was best captured by a model of a learner who is sensitive to pedagogy—that is, a learner who infers that the pedagogical demonstrator is a helpful teacher who is trying to show them how to get the reward, and so most likely performed the implausible action to demonstrate that it was causally necessary (Figure 1; Buchsbaum et al., under review). Thus, these findings offer a rational explanation for preschoolers' tendency to overimitate unnecessary actions.

Current evidence for whether toddlers reason about pedagogical contexts in the same way as preschoolers is mixed. For example, the inferences made by 3-year-olds about the generalizability of novel causal properties of objects (e.g., being magnetic) did not differ between intentional and pedagogical demonstrations (Butler & Markman, 2012), suggesting that pedagogical reasoning skills may develop across the preschool period. In contrast, other studies suggest that infants and toddlers *do* differentiate pedagogical demonstrations from non-social demonstrations (e.g., Brugger et al., 2007; Nielsen, 2006; Kupán et al., 2017; Shneidman et al., 2016), though recent work (Bazhydai et al., 2020) has failed to replicate an earlier finding that 2-year-olds preferentially transmit pedagogically demonstrated actions (Vredenburg, Kushnir & Casasalo, 2015).

None of these previous studies with toddlers have explicitly investigated the role of pedagogy in the imitation

of causal actions. It is possible that differences in pedagogical reasoning skills could explain the observed developmental increase in overimitative behavior across early childhood (e.g., McGuigan & Whiten, 2007; Chudek, Baron & Birch, 2016); perhaps toddlers are less likely than preschoolers to faithfully copy causally unnecessary actions because they do not make the same inferences about the demonstrator's pedagogical intent.

When might we expect pedagogical inference to emerge during development? The theory of 'natural pedagogy' (e.g., Csibra & Gergely, 2009) argues that humans possess an innate, unique adaptation for the social learning of causally-opaque actions. According to this account, from infancy, learners preferentially attend to information accompanied by ostensive-communicative cues (e.g., eye-contact, gaze alternation, learner-directed speech; e.g., Senju & Csibra, 2008; but see Gredebäck, Astor & Fawcett, 2018), and the pedagogical context fundamentally changes how a learner interprets the information they receive. On this account, we would expect toddlers to behave comparably to preschoolers in imitation tasks, because both age groups would be expected to privilege pedagogical cues in a social learning context (Bazhydai et al., 2020).

Alternatively, it is plausible that pedagogical inference requires the learner to explicitly represent the demonstrator's intent to teach. Indeed, the formal models discussed earlier assume that the learner represents that the evidence provided is being sampled by a knowledgeable, helpful, teacher. Thus, rather than being a result of an evolved, early-developing tendency to automatically privilege pedagogical cues, pedagogical inference might depend on cognitive abilities that continue to develop over the preschool period, such as theory of mind (e.g., Kline, 2015; Skerry et al., 2013; Shneidman & Woodward, 2015). In this case, toddlers might behave differently than preschoolers, given that on current evidence it is only by around 4 years of age that children have a robust understanding of the minds of others (e.g., Wellman, Corss & Watson, 2001; though note that there is ongoing controversy regarding whether there is evidence for 'implicit' belief reasoning in infants, e.g., Kulke et al., 2018).

The aim of the present study was to investigate which actions toddlers choose to copy following causal action demonstrations that differ in terms of the intentionality of the demonstrator. In particular, we were interested in whether toddlers, like preschoolers, would copy even causally implausible actions following a pedagogical demonstration. To address this question, we presented 18- to 30-month-olds with the same causal imitation task that Buchsbaum et al. (under review) presented to 3- to 5-year-old children.

Overview

In our Main Experiment, toddlers watched a demonstrator perform a sequence of two actions (AB) on a puzzle-box (e.g., spinning a dial and then pulling a lever; Figure 2) that led to a sticker being dispensed. We manipulated the intentionality of the demonstration (*unknowing*, *intentional*, *pedagogical*) between-subjects and the causal plausibility of

the first action (A) in the sequence (*connected*, *disconnected*) within-subjects. For a learner who is sensitive to pedagogical intent and causal plausibility, differentiation of the connected and disconnected conditions should be modulated by the demonstrator's intentionality. Specifically, there should be decreasing differentiation with increasingly explicit cues to intentionality, as predicted by a *pedagogically-sensitive* computational model (Figure 1) and seen in the behavior of preschoolers (Buchsbaum et al., under review). In terms of the behavioral data, this would be manifested as an interaction between causal plausibility and demonstrator intentionality, leading to comparatively faithful copying of the two-action sequence following a pedagogical demonstration, even in the disconnected condition.

We predicted that toddlers in the present study would be sensitive to the causal implausibility of a physically disconnected action leading to an effect (and therefore discriminate the connected and disconnected conditions), given that one-year-olds were sensitive to space and physical causality in an imitation paradigm (Brugger et al., 2007). We also expected toddlers to differentiate our unknowing and intentional demonstrations, given that within the first year of life, infants understand actions as goal-directed (e.g., Meltzoff, 1995; Woodward, 1998), and by 18 months toddlers distinguish accidental from intentional actions (Carpenter, Akhtar & Tomasello, 1998).

If toddlers automatically privilege pedagogical cues over other sources of information when deciding which actions to copy, as predicted by the theory of natural pedagogy (e.g., Csibra & Gergely, 2009), then like preschoolers, they should be more likely to copy the actions a pedagogical demonstrator performs, even those that are causally implausible (Buchsbaum et al., under review). If on the other hand there are developmental changes in pedagogical reasoning skills over this age range (e.g., because other cognitive abilities that continue to develop over the preschool period such as theory of mind are also necessary), then, unlike preschoolers, we would not expect toddlers to differentiate the intentional and pedagogical demonstrations.

Computational Model

In addition to collecting behavioral data from toddlers, we fit the data to a series of Bayesian computational models—an approach that enables us to formalize how learners combine the data they observe with their prior knowledge and expectations to make inferences (e.g., Perfors et al., 2011). Specifically, we fitted our data to models for learners that differed in their sensitivity to the demonstrator's intentionality (*intentionality blind*, *intentionality sensitive*, *pedagogically sensitive*; Figure 1), as well as their prior beliefs about the causal plausibility of disconnected actions (*connectedness* prior), and the likelihood of two- vs. single-action causes (*length* prior; see Buchsbaum et al., under review, for full details of the model). This allowed us to examine which model and combination of parameters best fit toddlers' performance in our task, and hence how these different factors influenced their behavior.

In this model, we assume that learners observe a demonstrator performing sequences of action, \mathbf{a} , which produce outcomes \mathbf{o} , and use Bayes' rule to evaluate the probability of a given hypothesis, h given this data:

$$P(h|\mathbf{a},\mathbf{o}) \propto P(\mathbf{o}|\mathbf{a},h)P(\mathbf{a}|h)P(h). \quad (1)$$

Here, hypotheses represent possible action sequences that could cause the puzzle-box to produce a reward. For instance, in our task, the hypothesis that, e.g., first the dial and then the lever is necessary corresponds to $h = AB$, while the belief that just the lever is necessary would correspond to $h = B$. We assume that learners are trying to bring about the desirable outcome (a sticker), and choose an action sequence to perform based on the probability that the sequence is causal. $P(\mathbf{o}|\mathbf{a},h)$ represents the probability of the outcome (e.g., sticker or no sticker) following the observed actions. If an action is effective, we assume that it will always cause the puzzle-box to produce a reward. For instance, if B is causal then both B and AB are effective action sequences.

The last term $P(h)$ represents the prior probability of each hypothesis, and is where we capture physical causal assumptions, for instance that actions on a physically separate box are less causally plausible. As a default, learners may assume a uniform prior over possible actions, meaning that they think all individual actions and sequences of actions are *a priori* equally likely. We also consider two sets of non-uniform prior biases that learners might have. The first is a *connectedness* prior where the probability that a causal sequence contains a disconnected action (e.g. action A on the physically separate box) is λ . Values of $\lambda < 0.5$ correspond to an increasingly strong belief that disconnected actions are unlikely to cause a reward (disfavoring any sequence containing A , when A is on the disconnected box). The second is a *length* prior where the likelihood that a causal action sequence contains just a single action is δ . Values of $\delta > 0.5$ correspond to an increasingly strong belief that multi-step causes are unlikely.

Finally, the critical piece is the middle term, $P(\mathbf{a}|h)$, which gives the probability that the demonstrator chose the observed set of actions to perform, given a specific true causal structure. This is where we can capture learners who interpret observed actions as being performed intentionally or pedagogically. In particular, knowing the demonstrator's goals may reduce uncertainty about the necessity of their actions. We assess three models of producing demonstrations; unknowingly, intentionally and pedagogically, and three kinds of learners who vary in whether they distinguish these demonstration types.

Following Buchsbaum et al. (under review), we model a demonstrator who brings about the outcome *unknowingly* as $p(\mathbf{a}|h) = 1/N$, where N is the number of possible actions the learner could perform, meaning that the demonstrator chose their actions randomly from the set of possible actions, without respect to their outcomes (Gweon, Tenenbaum, & Schulz, 2010; Xu & Tenenbaum, 2007). We model a demonstrator who acts *intentionally* in order to get the reward

as randomly sampling their actions from the set of actions that are effective at producing the desired outcome under the true causal structure, $P(\mathbf{a}|h) = 1/|E(h)|$. We model a demonstrator who performs each action *pedagogically*, as $P_{teacher}(\mathbf{a}|h) \propto P_{learner}(h|\mathbf{a},\mathbf{o})$, where $P_{learner}(h|\mathbf{a},\mathbf{o})$ is the learner's posterior probability of each hypothesis given the observed data (Bonawitz et al., 2011; Buchsbaum et al., 2011; Shafto et al., 2014).

Model Predictions

Figure 1 shows *a priori* model predictions for learners who are sensitive to causal plausibility, but differ in their sensitivity to the intentional evidence. An *intentionality-blind* learner does not differentiate between unknowing, intentional and pedagogical demonstrations, treating all demonstrations as unknowing. An *intentionality-sensitive* learner differentiates intentional actions from unknowing or accidental ones, but does not recognize pedagogical demonstrations as differing from intentional ones. Finally, a *pedagogically-sensitive* learner is able to distinguish unknowing, intentional and pedagogical demonstrations, and applies the appropriate model to each case. $\lambda = 0.25$ corresponds to a moderate preference for actions to be connected, and $\delta = 0.5$ corresponds to no pre-existing preference for shorter (or longer) sequences.




Figure 1: *A priori* model predictions for learners with different sensitivities to the demonstrator's intentionality: *intentionality-blind*; *intentionality-sensitive*; and *pedagogically-sensitive* ($\lambda = 0.25$ and $\delta = 0.5$)

Method

Study Design

All toddlers initially participated in the Baseline condition (one session) to establish their ability to learn to use the puzzle-box, followed by the Main Experiment (two sessions). There was a minimum gap of one week between sessions. The two sessions of the Main Experiment corresponded to our two within-subject causal plausibility conditions (*connected* and *disconnected*), which participants completed in a counterbalanced order. For the Main Experiment, participants were randomly assigned to one of three between-subject demonstrator intentionality conditions: *unknowing*, *intentional* or *pedagogical*.

Participants

Forty-two toddlers (22 males, $M_{\text{age}} = 23.98 \pm 0.46$ months, range = 18- to 30-months) completed the study. An additional 13 toddlers were tested but their data excluded for the following reasons: did not reach criterion of 5 puzzle-box activations in the Baseline condition (8); did not complete all three sessions of the study (4); caregiver interference (1).

Apparatus and Testing Setup

We used causal puzzle-boxes that dispensed stickers. The boxes had two distinct actions on the front (A and B, e.g., Figure 2) and contained remotely-operated sticker dispensers. Since an experimenter covertly triggered reward release when the participant manipulated the necessary action, either action could be ‘causal’. Toddlers were tested in a lab at the University of Toronto, or in an empty classroom at a local daycare.

Procedure

Two experimenters were involved in running the experiment: a demonstrator (D), and an experimenter (E) who accompanied the toddler. During each session, E and the child entered the testing area, to find D “busy” looking at a clipboard. E said “Hmmm, it looks like [D] is still getting ready, but we can wait here” and the child sat on a chair ~ 1 m from the apparatus. D performed a demonstration with the puzzle-box that resulted in a sticker being dispensed (the nature of the demonstration varied between the Baseline condition and the Main Experiment, and between the different demonstrator intentionality conditions; see below).

The demonstration was repeated until the subject had attended to two demonstrations. D then said “Oh hey, I’m all done! You can have a turn, and you can have these stickers!” D then stepped to the side of the room (Baseline condition), or left the room (Main Experiment), and E approached the puzzle-box with the child. If the child did not spontaneously interact with the puzzle-box, E provided neutral encouragement such as “It’s your turn, you can try anything!” Each time the child touched the causally necessary action a sticker was released, for up to 5 activations.

Baseline Session Without engaging in eye-contact with the participant, D performed a single action on the connected box, following which a sticker was dispensed.

Main Experiment D performed two actions, AB, on the box, following which a sticker was dispensed. Only B was causally necessary. The procedure varied according to the demonstrator’s (D) intentionality and the causal plausibility of the first action, A:

Unknowing: D faced sideways and pretended to read a clipboard as they performed the action sequence without looking at the participant or the puzzle-box. D did not see or acknowledge the sticker being dispensed from the box.

Intentional: D faced the participant but did not make eye-contact or acknowledge their presence. D looked at the puzzle-box as they performed the action sequence and saw the sticker being dispensed. They picked up the sticker and looked at it before putting it back down in front of the box.

Pedagogical: D faced the participant, engaged them in eye-contact, tapped on the top of the box(es) and said “Hey, [toddler name], look!” before performing the action sequence while alternating gaze between the participant and the box. D saw the sticker being dispensed, picked it up and said “Look, a sticker!” before putting it back down in front of the box.

Connected: Both actions in the sequence were performed on the puzzle-box that dispensed the sticker (Figure 2a).

Disconnected: The first action in the sequence (A) was performed on a physically disconnected box, separate from the puzzle-box that dispensed the sticker (Figure 2b).




Figure 2: Examples of the 2-action puzzle-boxes used in the (a) connected and (b) disconnected conditions.

Scoring and Analysis

Baseline condition We coded all touches of the demonstrated and non-demonstrated actions. To analyze toddlers’ tendency to act on the demonstrated action (A) versus the non-demonstrated action (B), we used a logistic repeated measures regression, with activation (1-5) as a covariate. We used binomial tests to analyze whether toddlers were more likely to perform the demonstrated action on their very first attempt to activate the box.

Main Experiment We coded all touches of the first (A) and second (B) actions in the sequence. For each activation of the puzzle-box, we scored whether it was preceded by at least one A touch (which reflects the extent to which the first, ambiguously necessary, action in the sequence (A) was incorporated). We analyzed subjects’ tendency to act on A first (as opposed to B), using a logistic repeated measures regression, with causal plausibility and intentionality as factors and activation (1-5) as a covariate.

Results and Discussion

Baseline session Toddlers were more likely to interact with the demonstrated action than the non-demonstrated action ($\beta_{\text{intercept}} = -4.41$, $s.e. = 1.11$, $p < 0.001$) and this did not change across activations ($\beta_{\text{activation}} = -0.46$, $s.e. = 0.28$, $p = 0.095$). They were also more likely to interact with the demonstrated action on their very first attempt to activate the box (37/42, $p < 0.001$). This provides evidence that toddlers learned how to activate the box by observing the demonstrator’s actions.

Main Experiment Analysis of Deviance (Type II Wald χ^2 test) on the logistic regression model showed that the extent to which toddlers first manipulated A when activating the puzzle-box was significantly influenced by causal plausibility ($\chi^2(1) = 17.25$, $p < 0.001$). Toddlers manipulated action A significantly more often when they saw a demonstration in which A was on the same box that produced

the reward (connected condition, Figure 2a; red bars in Figure 3a) compared with when action A was on a separate box (disconnected condition, Figure 2b; blue bars in Figure 3a). There was no effect of demonstrator intentionality ($\chi^2(2) = 2.38, p = 0.30$) and the interaction between intentionality and causal plausibility was not significant ($\chi^2(2) = 2.00, p = 0.37$). There was a significant effect of activation number ($\chi^2(2) = 20.29, p < 0.001$), with participants being less likely to manipulate A across activations (log odds ratio: 0.70:1).

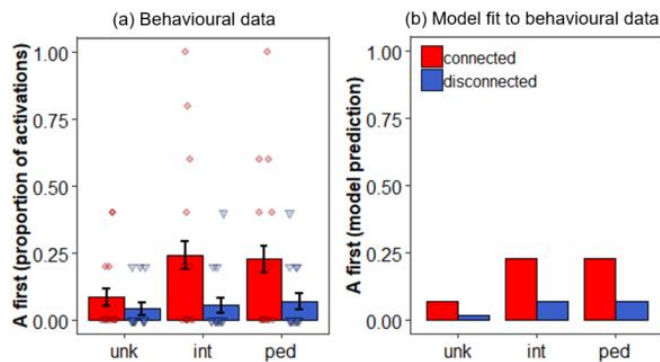


Figure 3: (a) Behavioural data (points are individual toddlers) and (b) best-fitting model results

Model fit The results of model-fitting provide additional evidence that, despite being sensitive to ostension, toddlers did not interpret an ostensive demonstration as indicative of pedagogical intent in this task. The *intentionality-sensitive* model favouring connected actions ($\lambda = 0.20$) and single-action causes ($\delta = 0.93$) gave the best fit to toddlers' behavioral data (Figure 3b). This model performed better than an *intentionality-blind* model ($\Delta AIC = 6.29$) and fit was not improved by assuming a *pedagogically-sensitive* learner ($\Delta AIC = 0.92$). The best-fitting model also performed better than an *intentionality-sensitive* model that assumed no preference for connected actions or shorter sequences ($\chi^2(2) = 691.64, p < 0.001, \Delta AIC = 687.64$), or that assumed a preference for connected action but not for shorter sequences ($\chi^2(1) = 262.16, p < 0.001, \Delta AIC = 260.16$).

These results suggest that although toddlers were sensitive to the causal plausibility of the first action in a sequence and so were less likely to perform A if it was on a disconnected box, their behavior was not strongly influenced by the demonstrator's intentionality. The best-fitting model suggests toddlers differentiated an intentional demonstration from an unknowing demonstration, which is in line with previous work with this age group (e.g., Carpenter et al., 1998; Moore et al., 2015). However, we found no evidence that toddlers discriminated actions performed intentionally from those performed pedagogically.

General Discussion

This study investigated how a demonstrator's intentionality and the causal plausibility of the first action (A) in a two-action sequence influenced toddlers' copying behavior. Of particular interest was whether toddlers, like preschoolers

tested previously (Buchsbaum et al., under review) would be more likely to copy an action sequence—even a causally implausible one—following a pedagogical demonstration, as opposed to a non-communicative demonstration. According to formal computational models, if a learner is sensitive to pedagogy they should be more likely to copy both actions, because of the assumption that the observed evidence is being provided by a knowledgeable, helpful, teacher, which leads to the inference that the actions are necessary.

According to our empirical data and the best-fitting model, toddlers were sensitive to the causal plausibility of the first action in the sequence: they were less likely to copy action A when it was on a physically disconnected box. A connectedness parameter (λ) of 0.20, which was similar to the value for preschoolers ($\lambda = 0.23$, Buchsbaum et al., under review), suggests that toddlers had a moderately strong belief that disconnected actions are an unlikely cause of a reward being dispensed. This is in line with previous research suggesting that even 1-year-olds are sensitive to contact relations between cause and effect (Brugger et al., 2007).

One concern might be that, rather than toddlers having any prior beliefs about the causal (im)plausibility of connected vs. disconnected actions in our setup, they were less likely to perform A in the disconnected condition due to it being further away from the sticker dispenser and thus more effortful to act on. However, 1- and 2-year-olds readily performed action A in the same disconnected puzzle-box setup when temporal information implied its causal necessity (Tecwyn et al., under review). This suggests that the additional distance/effort required to act on A relative to action B is unlikely to explain the overall pattern of omitting action A that was observed in the present study.

Toddlers in our study tended to omit action A across the conditions of our experiment (Figure 3a), copying it in just 12% of their puzzle-box activations overall. This stands in contrast to the behavior of 3- to 5-year-olds, who reproduced the two-action sequence on 67% of activations (Buchsbaum et al., under review). One potential explanation for toddlers' tendency to omit A is that they acted on the basis of a recency effect, which would lead them to act on the last action they saw the demonstrator interact with—in this case B. However, when 1- and 2-year-olds presented with the same puzzle-box as in the current study watched an adult perform an action (A), following which a sticker dispensed (effect E), following which a second action (B) was performed, they were significantly more likely to (correctly) manipulate A than B (Tecwyn et al., under review). This is the opposite of what would be predicted by a recency effect, so this is an unlikely explanation for the current data.

Another possibility is raised by the model-fitting results, which suggest that toddlers have a strong preference for single-action causes, as indicated by a length parameter (δ) of 0.93. This is in contrast to preschoolers, for whom including a length prior did not improve model fit to their data, suggesting that by four years of age children think single- and two-action causes are equally plausible (Buchsbaum et al., under review). If toddlers have a strong prior belief that two-

action causes are unlikely, then this could potentially be outweighing the influence of social cues in our task.

Prior work suggests that infants/toddlers may view 2-action sequences as unlikely causes. In a task where a 2-action sequence was causally necessary, only 29% of 15-month-olds reproduced the demonstrated sequence, compared to 39% who performed either the first or the second action, but not both (Brugger et al., 2007). Similarly, 14- and 18-month-old infants readily reproduced a single demonstrated action out of two possible actions (in line with the results of our Baseline condition), but only 6/20 infants spontaneously reproduced a demonstrated two-action sequence in the correct order (Carpenter et al., 1998). Although this previous work focused on younger children, ongoing work using the same causal puzzle-boxes and age group as the present study suggests that older toddlers may also struggle to learn causal structure when sequential actions are necessary (Tecwyn et al., 2020), and this possibility warrants further investigation.

The intentionality of the demonstrator did not have a significant effect on toddlers' copying behavior in our task, though the model-fitting results support the possibility that they may have differentiated unknowing demonstrations from intentional and pedagogical demonstrations. Importantly, toddlers were no more likely to copy the action sequence following a pedagogical demonstration compared to a non-communicative intentional demonstration. This suggests that, contrary to some prevalent accounts of pedagogy (e.g., Csibra & Gergely, 2009), the pedagogical demonstration did not result in a fundamental change in the interpretation of the observed information in our task.

It is possible that a pedagogical demonstration did not lead to faithful copying by toddlers because true pedagogical reasoning—where the learner represents the demonstrator's intent to teach (as specified in the computational model of a pedagogically-sensitive learner)—may be dependent on cognitive abilities that continue to develop into the preschool period. One potentially relevant cognitive skill is the ability to reason about others' mental states—or theory of mind—which is known to develop rapidly during the toddler and preschool years (e.g., Wellman et al., 2001). A starting point for investigating the relationship between pedagogical reasoning and understanding of mental states would be to investigate whether theory of mind skills predict toddlers' tendency to copy causally unnecessary actions performed by a pedagogical demonstrator.

Our results stand in contrast with some previous findings suggesting that toddlers *are* sensitive to social cues in imitation tasks. For example, 18- and 24-month-olds were more likely to copy a demonstrated method of achieving a goal when it was performed with accompanying ostensive cues, compared to when the demonstrator was non-communicative (Kupán et al., 2017; Nielsen, 2006), and 2-year-olds who saw a pedagogical demonstration of a toy's function spent less time exploring the rest of the toy than those who saw an intentional demonstration (Shneidman et al., 2016). Similarly, a study by Vredenburgh et al. (2015) suggested that 2-year-olds were more likely to transmit

pedagogically demonstrated actions to a naïve adult. However more recent work failed to replicate this effect (Bazhydai et al., 2020), concluding that toddlers were equally likely to transmit actions, regardless of the social context in which they were learned.

How might these previous findings be reconciled with the results of the current study? Notably, all of the studies described above only involved performing a single action as opposed to a two-action sequence, which may be crucial given that, as previously discussed, toddlers may struggle to grasp causal sequences, or at least find them highly implausible. More generally, between-study variations in the causal plausibility of the demonstrated actions might influence the role the pedagogical cues play—something that can be systematically investigated in future research.

In conclusion, unlike preschoolers, and contrary to what would be expected according to natural pedagogy accounts, toddlers were no more likely to faithfully copy an action sequence following a pedagogical demonstration vs. a non-communicative demonstration in a causal imitation task. This behavioral result was supported by formal computational modelling—toddlers' copying behavior was best fit by a model of a learner who is not sensitive pedagogy. These findings suggest that although preschoolers may infer that a demonstrator who uses ostensive cues is a knowledgeable, helpful teacher, toddlers may not make this same inference, at least in the present task. More work is needed to better understand the role of pedagogical cues in guiding learning at different stages of development, particularly regarding how social cues are weighted relative to other (non-social) cues in different contexts.

Acknowledgements

We thank Nina Esmail, Christine Heighton, Sarah MacKay Marton, Amanda Nickerson, Madeline Pelgrim and Pingki Mazumder for assistance with data collection. We are also grateful to participating daycares in Toronto. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC, 2016-05552).

References

- Bazhydai, M., Silverstein, P., Parise, E., & Westermann, G. (2020). Two-year-old children preferentially transmit simple actions but not pedagogically demonstrated actions. *Developmental Science*, e12941.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322-330.
- Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(2), 10918-10925.
- Brugger, A., Lariviere, L. A., Mumme, D. L., & Bushnell, E. W. (2007). Doing the right thing: Infants' selection of

- actions to imitate from observed event sequences. *Child Development*, 78(3), 806-824.
- Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120(3), 331-340.
- Buchsbaum, D., Tecwyn, E. C., Whalen, A., Messer, E. J. E., Bryant, E. L. F., Griffiths, T. L., Gopnik, A. & Seed, A. M. (under review). Children, but not capuchins, rationally integrate social and physical information when deciding which actions to copy.
- Butler, L. P., & Markman, E. M. (2012). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Development*, 83(4), 1416-1428.
- Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen-through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, 21(2), 315-330.
- Chudek, M., Baron, A. S., & Birch, S. (2016). Unselective overimitators: The evolutionary implications of children's indiscriminate copying of successful and prestigious models. *Child Development*, 87(3), 782-794.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148-153.
- Gredebäck, G., Astor, K., & Fawcett, C. (2018). Gaze following is not dependent on ostensive cues: a critical test of natural pedagogy. *Child Development*, 89(6), 2091-2098.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, 107(20), 9066-9071.
- Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Animal Cognition*, 8(3), 164-181.
- Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences*, 38, 1-71.
- Kulke, L., von Duhn, B., Schneider, D., & Rakoczy, H. (2018). Is implicit theory of mind a real and robust phenomenon? Results from a systematic replication study. *Psychological Science*, 29(6), 888-900.
- Kupán, K., Király, I., Kupán, K., Krekó, K., Miklósi, Á., & Topál, J. (2017). Interacting effect of two social factors on 18-month-old infants' imitative behavior: Communicative cues and demonstrator presence. *Journal of Experimental Child Psychology*, 161, 186-194.
- Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences*, 104(50), 19751-19756.
- Meltzoff, A. N. (1995). Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5), 838-850.
- Moore, R., Mueller, B., Kaminski, J., & Tomasello, M. (2015). Two-year-old children but not domestic dogs understand communicative intentions without language, gestures, or gaze. *Developmental Science*, 18(2), 232-242.
- Nielsen, M. (2006). Copying actions and copying outcomes: social learning through the second year. *Developmental Psychology*, 42(3), 555-565.
- Perfors, A., Tenenbaum, J. B., Griffiths, T. L., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120(3), 302-321.
- Senju, A., & Csibra, G. (2008). Gaze following in human infants depends on communicative signals. *Current Biology*, 18(9), 668-671.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55-89.
- Shneidman, L., Gaskins, S., & Woodward, A. (2016). Child-directed teaching and social learning at 18 months of age: evidence from Yucatec Mayan and US infants. *Developmental Science*, 19(3), 372-381.
- Shneidman, L., Gweon, H., Schulz, L. E., & Woodward, A. L. (2016). Learning from others and spontaneous exploration: A cross-cultural investigation. *Child Development*, 87(3), 723-735.
- Skerry, A. E., Lambert, E., Powell, L. J., & McAuliffe, K. (2013). The origins of pedagogy: Developmental and evolutionary perspectives. *Evolutionary Psychology*, 11(3), 147470491301100306.
- Tecwyn, E. C., Mazumder, P., & Buchsbaum, D. (under review) Toddlers' causal interventions reveal a grasp of temporal priority.
- Tecwyn, E. C., Mahbub, N., Kazi, N., & Buchsbaum, D. (2020). Can Toddlers Learn Causal Action Sequences? *Poster at Cognitive Science Society Meeting*.
- Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology*, 28(1), 509-529.
- Vredenburg, C., Kushnir, T., & Casasola, M. (2015). Pedagogical cues encourage toddlers' transmission of recently demonstrated functions to unfamiliar adults. *Developmental Science*, 18(4), 645-654.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655-684.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1-34.
- Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, 10(3), 288-297.