

Children’s causal inferences about past vs. future events

Katharine A. Tillman
Department of Psychology
The University of Texas at Austin
ktillman@utexas.edu

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

Abstract

Causal and temporal reasoning are fundamentally linked, but few studies have directly examined how the ability to make causal inferences about the past vs. the future develops. We used a counterfactual reasoning task to explore 4- to 6-year-old children’s understanding of the causal relationships among past, present, and future events. Like adults, even 4-year-olds judged that future, but not past, events could be altered by interventions in the present. This early sensitivity to the causal asymmetry between the past and future became more pronounced with age. We also found that children and adults selectively and appropriately use evidence about the present to make inferences about past events. Implications for theoretical accounts of the development of causal reasoning and abstract concepts of time are discussed.

Keywords: cognitive development; temporal cognition; causal inference; counterfactual reasoning

Introduction

You can change the future, but you can’t change the past. This fundamental distinction between the past and future is central to an abstract, linear concept of time, and has profound effects on adults’ everyday behavior. Although philosophers and physicists have argued about the ultimate reality of the past/future asymmetry, many of us find it difficult, if not impossible, to conceive of a world without it. Is the past/future distinction a “built in” feature of human cognition? If it isn’t, when and how does it develop? While we know that children’s reasoning about both temporal and causal relationships improves during the preschool years, few studies have directly explored the relationship between children’s reasoning about causality and their knowledge of the ontological distinction between past and future (see McCormack & Hoerl, 2017). Here, we use a counterfactual¹ reasoning task to explore how children use information about present events to make causal inferences about the past and future.

Despite the central role of the past/future distinction in explicating temporal reasoning, it remains unclear to what extent young children possess this understanding. While it is difficult to test this in preverbal infants, researchers have

looked to children’s earliest production of temporal language for clues. The past-tense verb marking *-ed* is one of the first grammatical inflections English-speaking children produce, usually at or before age 2 (Brown, 1973), which has been taken as evidence that understanding of the past/future status of events relative to the present develops early. However, there is debate in the language acquisition literature over how accurate and generalizable children’s early uses of tense are, and particularly whether they may indicate perfective aspect, rather than event time (Anderson & Shirai, 1996). Deictic time words like “tomorrow” and “yesterday” are also early to appear in the child’s lexicon, though children don’t use them reliably for several years (e.g., Tillman et al, 2017). Nonetheless, while these studies suggest early onset and prolonged development of past-future reasoning, it remains possible that children’s understanding the causal asymmetry of the past and the future develops prior to the ability to express these differences in language.

When do children understand how causality operates over time? Suggesting that even infants intuitively understand the relationship between temporal order and causality, 4-month-olds look longer when presented with impossible causal chains of events, including those with apparent breaks in temporal continuity (Cohen et al., 1998). Nevertheless, recognizing the temporal-causal structure of a simple event, like one ball striking another, does not imply that infants have a concept of the past or the future. Later in development, when asked which of two possible events caused another event to happen, 3-year-olds chose the prior rather than the subsequent one (Bullock and Gelman, 1979). Four-year-olds recognize that past, but not future, events determine present mental states and states of the world (Busby Grant & Suddendorf, 2010). However, 4-year-olds struggle to use information about the relative ordering of multiple past events to make inferences about the present, and fail to solve temporal reasoning tasks in which the order they receive information about events doesn’t match the order in which those events occurred (e.g., McCormack & Hoerl, 2007), suggesting that young children lack flexible temporal perspective-taking.

¹The precise definition of counterfactual reasoning, and thus the age at which children are first capable of it, is the subject of much debate (e.g., Beck, 2016; Weisberg & Gopnik, 2013). In the present paper, we take a broad view of

counterfactual reasoning, which encompasses hypothetical questions about the past, present, and future, as well as conditionals.

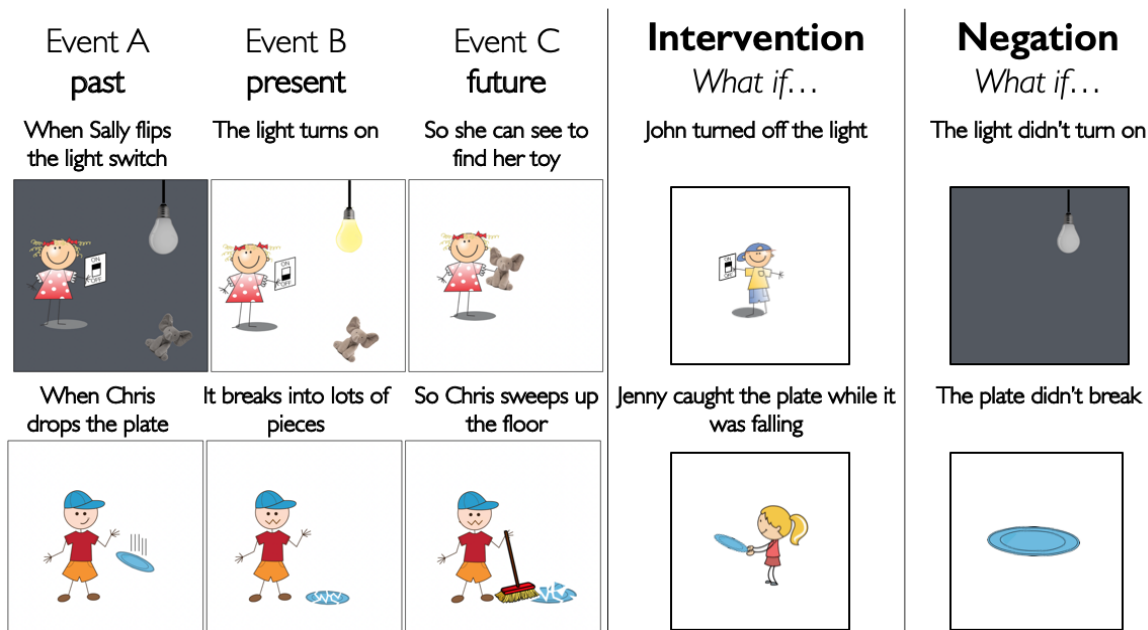


Figure 1: Example storyboards and intervention/negation cards. The experimenter (E) recited the story on the front of the card, then flipped it to reveal three empty boxes. At test, E placed either an intervention or negation card in the center square (B) before posing counterfactual questions about its effects on past (A) and future (C) events.

Critically, none of these prior studies directly test whether children understand that the future is alterable in the present, but the past is not. Instead, a separate literature has explored the development of children’s causal reasoning skills (see Gopnik & Wellman, 2012 for a review), and, despite the related subject matter, this literature has developed largely independently from the work on the development of temporal cognition discussed above. Within the causal reasoning literature, some accounts suggest that causal relationships are defined in terms of their counterfactual dependency. That is, if event A causes event B, then an intervention on A will lead to a change in B (Pearl, 2000). Relevant to this, researchers have examined children’s understanding of this link between causal and counterfactual reasoning. For example, in early work, Harris, German, and Mills (1996) conducted a series of studies in which 3- and 4-year-olds were presented with scenarios in which they were asked to reason about a short causal sequence (e.g., a character walks across the floor with muddy boots [A], making a mess [B]). When asked conditional questions about what would have happened had A not occurred (e.g., “What if Carol had taken her shoes off? Would the floor be dirty?”), children made accurate judgements about effects on B.

If children have a unidirectional view of causality, given a 3-step causal chain of events, $A \rightarrow B \rightarrow C$, they should judge that an intervention at B can alter future event C, as previous studies have found. Importantly, however, they should also judge that the intervention will *not* alter past event A (Sloman & Lagnado, 2005).

Consider the following scenario (Figure 1): *When Sally flips the lightswitch, then the light turns on, so she can see to find her toy.* If told that another character, John, turned off the light² (at B), adults may reasonably predict that Sally will no longer be able to see (at C). However, they should not infer that Sally never turned the light on in the first place (at A), because John’s actions at time B can’t change what happened in the past, at time A. Here we test whether 4- to 6-year-old children make the same inferences.

Retrospective reasoning

Despite understanding time itself to be linear, under some circumstances, adults use information about the present to reason “back in time” and make inferences about what already happened in the past. For example, if an expected event does not occur—and no other explanation for this is given—an adult might reasonably infer that the event’s usual cause must not have occurred. For example, when simply told

² Note that under some definitions of counterfactual reasoning, our use of the simple past tense here, rather than pluperfect subjunctives such as “What if John *had* turned off the light?” or (in the negation condition) “What if the light *hadn’t* turned on?”, indicates that these statements are

hypothetical rather than counterfactual *per se* (see Lucas & Kemp, 2015). We chose simpler language primarily to make the task more comprehensible to young children, but are currently exploring whether this tense modification impacts performance on our task.

that the light *didn't* turn on (at B; see footnote 2), an adult might indeed conclude that Sally never flipped the lightswitch (at A). Here, we presented children and adults with stories involving 3-step causal chains, and then asked them to consider scenarios in which the second step (B) was different. Importantly, we asked both about the effects of the “present” change on the future (C) and on the past (A).

Given prior work showing that adults generate different causal predictions following passive observation (e.g., observing that B did not occur) than they do following interventions (e.g. acting on B to prevent it from occurring; Sloman & Lagnado, 2005; Waldmann and Hagmeyer, 2005), we also varied this feature in the current study. In the *intervention* condition, an external agent (e.g., another character) caused the “present” change. We hypothesized that participants with a linear concept of time would judge that the future event would also change, but not the past event. In the *negation* condition, however, no explanation for the present change was given. Here we hypothesized, again, that participants with a linear concept of time would judge that the future would change. If participants also engage in retrospective causal reasoning, we hypothesized that, unlike in the intervention condition, they would also systematically judge that the past event (A) had changed (e.g., Sloman & Lagnado, 2005). In contrast, if participants do not reason retrospectively, we predicted that they would perform similarly in the two conditions.

Method

Participants

A total of 258 subjects participated, including 65 4-year-olds ($M_{\text{age}} = 4.5$ years, range = 4.0-5.0 years), 70 5-year-olds ($M_{\text{age}} = 5.5$ years, range = 5.0-6.0 years), 63 6-year-olds ($M_{\text{age}} = 6.4$ years, range = 6.0-7.0 years) and 60 adult controls ($M_{\text{age}} = 21.6$ years, range = 18.2-31.1 years). Participants were pseudo-randomly assigned to either the Intervention or Negation condition. An additional 43 children participated, but were not included in analyses due to being outside the target age range ($n = 3$), experimenter or technical error ($n = 5$), failure to complete the task ($n = 4$), developmental delay ($n = 2$), insufficient fluency in English ($n = 1$), incomplete age information ($n = 2$), or failing more than one control trial, as described below ($n = 26$).

Materials

Study materials included eight 3-panel storyboards illustrating sequences of events from left-to-right. Two examples are shown in Fig 1. Each panel was 2.8 in. \times 2.8 in. Single images corresponding to event B in each story were also used in testing, which represented either identical pictures (control stories), interventions, or negations, depending on condition. Each individual image was square with a black outline, and on the reverse side of each storyboard were three empty black squares positioned like the filled images on the front of the card.

Procedure

Children were tested one-on-one, in a quiet room with the experimenter. The experimenter began the session by placing the first storyboard in front of the participant, saying “I’m going to tell you some stories. There are three things that happen in each story, see?” She then pointed to each image in the story while reciting the corresponding part of the narrative, in this case, “When [A] Julie opens the door, then [B] her dog runs outside, so [C] he smashes up all the flowers in the garden”.

The experimenter then flipped over the storyboard, revealing the 3 empty boxes, and initiated a demonstration control trial. While placing a duplicate of the center image from the front of the card in the empty center square, she asked the child “tell me, *just like in that story*, if [the dog ran outside]...,” and then pointed to the empty third [event C; future] box while completing the question with a forced choice: “will [he smash all the flowers in the garden] or not [smash the flowers in the garden]?” After receiving a verbal response from the participant, the experimenter repeated the procedure again, instead pointing to the first [event A; past] box and asking, “did [Julie open the door] or not [open the door]?”

Next, the experimenter flipped the card back over, repeated the original story, and explained that she would now be asking the participant to think about what would happen in the story if something had been *different*. In this demonstration critical trial, the experimenter placed a modified image B in the empty center square on the back of the card. This image showed either the **intervention**, “What if the dog were on a leash and couldn’t get out?”, or the **negation**, “What if the dog *didn't* run outside?”, according to condition. The test concluded with the past and future test questions, as above. No feedback was given on either the control or critical trials using the demonstration story.

After this demonstration phase, the experimenter told the participant that she would tell some other stories, sometimes asking if things had been the same (control stories), and sometimes asking if things had been different (critical stories), and sometimes asking about the first part or the story, and sometimes about the last part. The remainder of the task included 7 new stories, 5 of which were used on critical trials (see examples in Fig. 1) and 2 were control stories. The ordering of the stories (other than the demonstration story), the past and future questions about each story, and the positive and negative response options in each question were counterbalanced across participants. The third and sixth stories were always control stories. Procedures used in the intervention and negation conditions were identical, apart from the different counterfactual questions and corresponding images used during test.

Data from children who responded incorrectly more than one control trials, i.e., by denying that an event from the story they had just heard had occurred in that story, were excluded due to suspected incomprehension of the task. This exclusion criterion was particularly important because the predicted “adult-like” response pattern in the negation condition was

one in which the participant judges that *none* of the events in either past or future critical trials had occurred. We therefore wanted to minimize the chances of potentially confusing a “no bias” in children who did not comprehend the task at all with adult-like conditional reasoning.

Coding

During testing, the experimenter recorded whether the participant affirmed or denied that each past or future event would occur. Yes responses were coded as 1, no responses as 0. These were later reverse-coded as described below. Participants who answered more than one of the four control questions incorrectly (i.e., by responding that the event did not occur; $n = 26$) were excluded from further analysis. Data from the demonstration story were not included in analysis. All analyses were conducted in R, using the *lme4* package for mixed-effects modeling.

Results

We began with two primary questions about our dataset: (1) Do participants differentiate past from future in the intervention condition? and (2) Do participants ever reason retrospectively (i.e., “back in time”) when answering questions about the past in the negation condition?

Before addressing these, we asked whether children’s performance differed between the two conditions. Because our DV was a binary choice (either an event would occur or not), we conducted a mixed-effects logistic regression. For ease of exposition, the data were reverse-coded, such that answers indicating that events would *not* occur in the counterfactual scenario were considered “changes” (1), while answers indicating that events would still occur were considered non-changes (0). We modeled the likelihood that a child³ would say an event changed as a function of their age (continuous; between-subjects), condition (intervention vs. negation; between-subjects), and event time (past vs. future; within-subjects). We also included an interaction between event time and condition in this model, and random intercepts for subjects and stories. Results of this analysis revealed significant main effects of age ($\beta = 0.5, p = 0.004$) and event time ($\beta = -3.4, p < 0.001$) as well as a significant interaction between event time and condition ($\beta = 2.3, p < 0.001$). Given the evidence that children’s behavior differed between conditions, we proceeded to analyze the data from the two conditions separately.

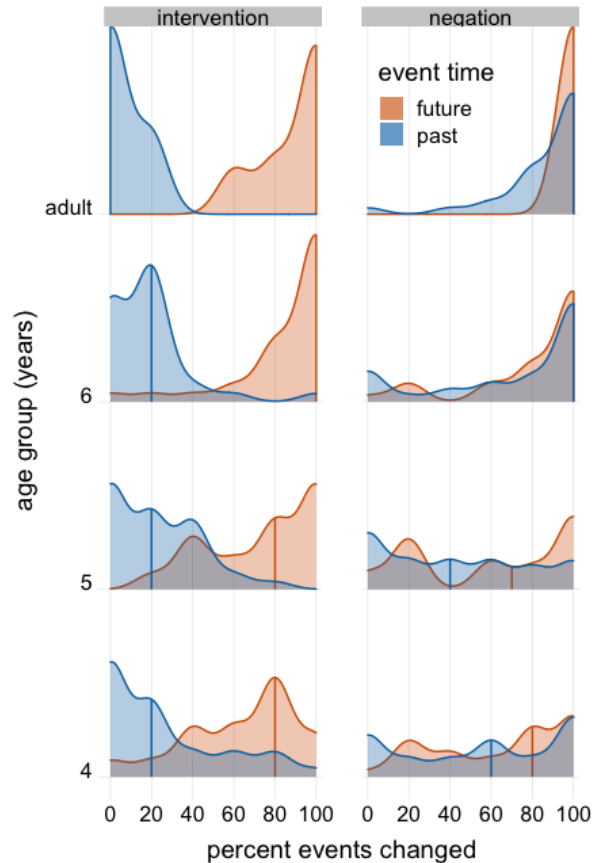


Figure 2: Distributions of responses to past (blue) and future (red) counterfactual questions, in the intervention (left) and negation (right) conditions. Height of shaded areas indicates the density of responses at each level of consistency, e.g., 80% = 4 of 5 events changed. Vertical lines = medians. Density calculation bandwidth = 8.

Intervention condition

Our goal in the intervention condition was to test whether participants differentiate the effects of present interventions on past vs. future events. In other words, do they know that you can change the future, but not the past? The distributions of responses to past- and future- questions, i.e., the percentage of target events that changed, for each age group are shown in the left column of Figure 2, with medians represented by vertical lines. As expected, and in line with prior work, adults strongly distinguished past from future: the median percentage of past events they said would change as a result of the intervention was 0%, 95% CI [0%-0%], while the median percentage of future events that would change was 100% [80%-100%].

³ Adult controls were not included in this analysis.




Figure 3: Proportion of subjects in each age group who demonstrated each of 4 response patterns across stories, in the intervention (top) and negation (bottom) conditions. “Linear” (green) = future events consistently judged to change, past events to stay the same. “Both change” (red) = both past and future events judged to change. No-changes (blue) = both past and future events were judged to stay the same. Mixed (purple) = inconsistent responses.

Considering the developmental data, a logistic regression of the children’s likelihood of saying that an event changed revealed significant effects of age ($\beta = 0.7, p < 0.001$) and event time ($\beta = 2.5, p = 0.03$) as well as a significant interaction ($\beta = -1.1, p < 0.001$). As shown in Fig. 2, children were more likely to judge that interventions would change the future than the past, and this effect increased with age. Wilcoxon signed-rank tests confirmed that the past vs. future effect was significant even in 4-year-olds, who reported that 80% [60%-80%] of future events changed, but only 20% [0%-20%] of past events did ($W = 400, p < 0.001$).

Interestingly, the three groups of children did not differ in their likelihood of judging that past events changed (Wilcoxon rank-sum tests, 4’s vs 6’s, $W = 567, p = 0.6$), though even 6-year-olds were significantly more likely to do so than were adults ($W = 306, p = 0.006$). On future questions, 4- and 5-year-olds were significantly less likely to say that events changed than were 6-year-olds and adults (5’s vs 6’s, $W = 731, p = 0.04$), though neither of these pairs differed (4’s vs 5’s $W = 456, p = 0.09$; 6’s vs adults $W = 489, p = 0.9$).

In addition to overall performance on past vs future questions, we were interested in the patterns of responses provided by individual subjects. For instance, did children who said the past wouldn’t change also say the future *would* change, as a linear model of time would predict? For the purpose of this analysis, we operationalized a “linear” pattern as one in which the participant judged that at least 4 out of 5 future events would change after intervention, *and* that at least 4 of the 5 past events would not. As shown in Figure 3 (top panel, green line), we found that 83% [65%-94%] of

adults conformed to this pattern, as did 72% [53%-86%] of 6-year-olds, 36% [21%-54%] of 5-year-olds, and 27% [13%-46%] of 4-year-olds. Subjects who didn’t follow a linear pattern typically reported fewer than 4 changes to future questions, resulting in a mixed pattern. Patterns in which either both events or neither event changed were rare in all age groups.

Negation condition

In the negation condition we assessed whether participants would reason retrospectively, making the inference that an observed, unexplained change in the present was caused by a prior change in the past. As shown in Fig 2 (right column), we found strong retrospective reasoning in adults: when simply told that the present event “didn’t” occur, they judged that the future effect would not occur on a median of 100%, 95% CI [100%-100%], of trials, and, in contrast to the intervention condition, that the past had changed on 100% [80%-100%] of trials, in line with prior adult work (e.g., Waldmann & Hagmeyer, 2005).

Next we considered the developmental data. A logistic regression model of the children’s data in the negation condition, with the same effects structure as the one used in the intervention condition, revealed only a main effect of age ($\beta = 1.1, p = 0.04$). Older children were more likely than younger children to judge that events changed. However, unlike in the intervention condition, there was no significant effect of event time ($\beta = 0.35, p = 0.8$), and no interaction ($\beta = -0.29, p = 0.24$). In other words, we did not detect evidence

that children were treating past and future events differently in this condition. Examining the age effects further, we found that 5-year-olds were less likely to judge that past events had changed than 6-year-olds and adults (5's vs 6's, $W = 302$, $p = 0.002$), but no other age-group comparisons reached significance. On future questions, 4- and 5-year-olds could not be distinguished, but 6-year-olds were significantly more likely to say that future events changed than were 5-year-olds ($W = 677$, $p = 0.04$), and adults were more likely to do so than were 6-year-olds ($W = 660$, $p < 0.001$).

Importantly, retrospective reasoning about the implications of the present on the past, when combined with knowledge of how the present influences the future (i.e., future “prospective” reasoning), predicts not only that the past and future will *not* be differentiated, as we found above, but also that both will be judged to have changed. This pattern was less common across the 4- and 5-year-old groups, as can be seen in the flatter distributions (broader confidence intervals) in Fig 2. For example, the median percentage of past events that 4-year-olds judged to have changed was 60% [40%-80%], and for 5-year-olds was 40% [20%-60%]. In contrast, the median percentage for both 6-year-olds and adults was 100%. In our individual-subjects analysis, a consistent retrospective/prospective reasoning pattern was operationalized as one in which at least 4 or 5 past events and 4 of 5 future events changed. As shown in Fig 3, we found that 41% [24%-59%] of 4-year-olds, 24% [11%-41%] of 5-year-olds, 65% [45%-81%] of 6-year-olds, and 87% [69%-96%] of adults displayed this pattern, while linear response patterns were very rare in this condition. Interestingly, among children who did not show this pattern, particularly 5-year-olds (who were surprisingly less adult-like than 4-year-olds), a larger proportion said that *neither* event changed (Fig. 3, blue lines) than we observed in the intervention condition. We discuss this further below.

Discussion

In the current study, we explored the development of children's reasoning about causal relationships among events in the past, present, and future. We discovered that children as young as 4 already distinguish the past and future: they are more likely to judge that an intervention in the present will change a future event than a past one. To our knowledge, this is the strongest evidence to date that pre-school children appreciate the causal asymmetry between the past and future (see McCormack & Hoerl, 2017). Moreover, children treated counterfactual scenarios with an explicit causal agent differently from those in which the cause must be inferred. In the latter case, children did not distinguish past from future, and by age 6, 65% of children consistently demonstrated both prospective and retrospective causal reasoning.

To test their reasoning about past and future events, we told participants 3-step stories, and then asked them to consider counterfactual cases in which an outside agent disrupted the middle step. We found that even 4-year-olds very rarely judged that the past event would retroactively change. This finding extends previous literature showing that 4-year-olds

understand that past (but not future) events can cause present ones (e.g., Busby & Suddendorf, 2010), and suggests that the understanding that time is irreversible is strong and early-developing. However, despite the high overall rate of denials that the past would change, we only found a consistently “linear” response pattern in about a quarter of 4-year-olds. This was because, compared to older children, younger children were less likely to judge that *future* events would change after intervention.

Finding more adult-like behavior from children on past than future trials is somewhat surprising in light of previous studies showing that 3-year-olds are capable of prospective (“forward”) conditional reasoning (e.g., Harris et al., 1996). In fact, it has been proposed that future conditionals are easier than past counterfactuals for children, because they do not require them to hold both the real world, i.e., how things actually occurred, and the possible world, i.e., how things could have been otherwise, in mind simultaneously (e.g., Beck, 2016; Beck & Riggs, 2014; Raefsteder et al. 2010). Perhaps, however, children in our task were more variable in their predictions about the future than the past simply because the future is intrinsically more open-ended. In linear time, a given intervention may or may not be effective at generating a particular outcome, but will *never* change what has already occurred.

Although children consistently denied that the past would change in the intervention condition, their judgments about past events were not rigid. In the negation condition, children's responses to past questions were more mixed, and like adults, they were more likely to say that the past *had* changed than that it hadn't. Given the minimal changes to the task across conditions, our finding that children are already sensitive to the precise nature of the conditional statement, and to the increased ambiguity of negations relative to explicit interventions (with respect to the past) is striking. Given children's high performance in the intervention condition, and the lack of a bell curve centered around random responding in the negation condition, we do not believe these results can be attributed to confusion about the nature of the task. Instead, these findings may suggest that some children (but not others) are already able to reason backward in time.

One intriguing possibility is that children who perform like adults in the negation condition are deploying what the adult counterfactual reasoning literature has termed “backtracking” (i.e., engaging in a special type of counterfactual reasoning that involves inference about upstream causal variables; Gerstenberg, Bechlivanidis, & Lagnado, 2013; Rips, 2010; Rips & Edwards, 2013; Sloman & Lagnado, 2005). Although there has been substantial debate over what types of counterfactuals lead to backtracking inferences in adults (e.g., Han et al., 2014), these investigations have not yet been extended to children. The tendency to engage in backtracking (or not) has important implications for interventionist accounts for causal reasoning. Specifically, because evaluating the effects of an intervention on a given variable requires “cutting off” that

variable from its upstream causal antecedents, backtracking should not be possible (Pearl, 2000; see Lucas & Kemp, 2015). While our negation condition is similar to certain backtracking tasks previously used with adults (e.g., Han et al., 2014; Lucas & Kemp, 2015), given the methodological differences (e.g., using child-friendly events that operate over time; presenting interventions/negations in past vs. pluperfect tense), additional work will be required to explore this potential developmental link.

In sum, the current study brings together the literatures on the development of causal reasoning and temporal cognition, by leveraging a counterfactual reasoning task to explore children's understanding of the past and the future. We found that children are able to recognize the causal asymmetry between past and future prior to the age of 4, reflecting the early development of a linear view of time.

Interestingly, it has been hypothesized that counterfactual reasoning itself may hinge on the development of an abstract, event-independent concept of time (McCormack & Hoerl, 2017). To consider different possible worlds, one must separate the time-point at which an event occurred from the event itself. Linear time thus provides a framework in which events can be organized and even mentally "switched out," so that their causal consequences can be considered. By studying these phenomena in tandem, future studies may uncover new insights about how time and causality are mentally represented.

References

- Anderson, S. W. & Shirai, Y. (1996). The primacy of aspect in first and second language acquisition: the pidgin-creole connection. In W. C. Ritchie & T. K. Bhatia (eds), *Handbook of Second Language Acquisition*. London: Academic Press.
- Beck, S. R. (2016). Counterfactuals matter: A reply to Weisberg & Gopnik. *Cognitive science*, 40(1), 260-261.
- Beck, S. R., & Riggs, K. J. (2014). Developing thoughts about what might have been. *Child development perspectives*, 8(3), 175-179.
- Brown, R. (1973). *A first language: The early stages*. Harvard U. Press.
- Bullock, M., & Gelman, R. (1979). Preschool children's assumptions about cause and effect: Temporal ordering. *Child Development*, 89-96.
- Busby, J. G., & Suddendorf, T. (2010). Young children's ability to distinguish past and future changes in physical and mental states. *British Journal of Developmental Psychology*, 28(4), 853-870.
- Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant behavior and development*, 21(4), 713-731.
- Gerstenberg, T., Bechlivanidis, C., & Lagnado, D. A. (2013). Back on track: Backtracking in counterfactual reasoning. In M. P. Knauff, M., N. Sebanz & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6), 1085.
- Han, J. H., Jimenez-Leal, W., & Sloman, S. (2014). Conditions for backtracking with counterfactual conditionals. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 36, No. 36).
- Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, 61(3), 233-259.
- Lucas, C. G. & Kemp, C. (2015). An improved probabilistic account of counterfactual reasoning. *Psychological Review*, 122(4), 700-734.
- McCormack, T., & Hoerl, C. (2007). Young children's reasoning about the order of past events. *Journal of experimental child psychology*, 98(3), 168-183.
- McCormack, T., & Hoerl, C. (2017). The development of temporal concepts: Learning to locate events in time. *Timing & Time Perception*, 5(3-4), 297-327.
- Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge University Press, Cambridge, UK.
- Rafetseder, E., Cristi-Vargas, R., & Perner, J. (2010). Counterfactual reasoning: Developing a sense of "nearest possible world". *Child development*, 81(1), 376-389.
- Rips, L. J. (2010). Two Causal Theories of Counterfactual Conditionals. *Cogn. Sci.*, 34(2), 175-221.
- Sloman SA, Lagnado DA. (2005). Do we "do"? *Cogn. Sci.* 29(1):5-39
- Rips, L. J., & Edwards, B. J. (2013). Inference and explanation in counterfactual reasoning. *Cognitive Science*, 37(6), 1107-1135.
- Sloman, S., & Lagnado, D. A. (2005). Do people 'do'. *Cogn. Sci.*, 29, 5-39.
- Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual review of psychology*, 66, 223-247.
- Tillman, K. A., Marghetis, T, Barner, D., & Srinivasan, M. (2017). Today is tomorrow's yesterday: Children's acquisition of deictic time words. *Cognitive Psychology*, 92, 87-100.
- Waldmann, M. R. , & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge . *Journal of Experimental Psychology: Learning, Motivation, and Cognition*, 31, 216-227.
- Weisberg, D. S., & Gopnik, A. (2013). Pretense, counterfactuals, and Bayesian causal models: Why what is not real really matters. *Cognitive science*, 37(7), 1368-1381.