Running head: CAUSAL STRUCTURE AND INTERVENTION JUDGMENTS

Key words: causation, causal reasoning, counterfactual reasoning, development

Are causal structure and intervention judgments inextricably linked? A developmental study.

Caren A. Frosch

Teresa McCormack

School of Psychology

Queen's University Belfast

David A. Lagnado

Division of Psychology and Language Sciences

University College London

Patrick Burns

School of Psychology

University of Birmingham

Address for Correspondence:

Teresa McCormack

School of Psychology

Queen's University Belfast

Belfast, Northern Ireland

BT7 1NN

t.mccormack@qub.ac.uk

Telephone: ++44 (0)2890 974174

Fax: ++44 (0)2890 975486

Abstract

The application of the formal framework of causal Bayesian Networks to children's causal learning provides the motivation to examine the link between judgments about the causal structure of a system, and the ability to make inferences about interventions on components of the system. Three experiments examined whether children are able to make correct inferences about interventions on different causal structures. The first two experiments examined whether children's causal structure and intervention judgments were consistent with one another. In Experiment 1, children aged between 4 and 8 years made causal structure judgments on a three component causal system followed by counterfactual intervention judgments. In Experiment 2, children's causal structure judgments were followed by intervention judgments phrased as future hypotheticals. In Experiment 3, we explicitly told children what the correct causal structure was and asked them to make intervention judgments. The results of the three experiments suggest that the representations that support causal structure judgments do not easily support simple judgments about interventions in children. We discuss our findings in light of strong interventionist claims that the two types of judgments should be closely linked.

Are causal structure and intervention judgments inextricably linked? A developmental study.

Causal knowledge is crucial to our ability to predict, control and explain the world around us. How do people acquire this knowledge in the absence of explicit instruction? Most previous research has focused on how the strength of causal relations is estimated from patterns of covariation (e.g., Cheng, 1997; Shanks, 2004). Thus in a typical laboratory-based causal learning task, participants are told which events are causes and which are effects, and have to judge the *causal strength* of a proposed relation. However, in the real world events are not always pre-sorted into causes and effects. We often have to make inferences about the *causal structure* of event relations. For example, suppose one observes that events A, B, and C are highly correlated. Do the events form a causal chain ABC, a causal chain ACB, or does A independently cause both B and C (see Figure 1)? Studies of causal structure learning examine how we decide between such possibilities, and recently such learning has attracted considerable research interest (Gopnik & Schulz, 2007; Kushnir, Gopnik, Lucas, & Schulz, 2010; Lagando & Sloman, 2004, 2006; Sobel & Kushnir, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). Young children in particular are likely to encounter situations in which they need to figure out causal structure, because they may often lack the substantive knowledge that would fully specify a model of the relations between variables (Gopnik et al., 2004).

Three recent studies have suggested that young children can discriminate between causal structures either on the basis of conditional probability information (Schulz, Gopnik, & Glymour, 2007; Sobel & Sommerville, 2009), or on the basis of temporal cues (Burns & McCormack, 2009). The first two experiments in the current paper follow up this research by also examining children's causal structure learning. However, the core aim of this paper is not to demonstrate that children can discriminate between different causal structures, but to test claims about the sorts of judgments children can make using the causal structures they have learned. Recent claims about how to model such learning suggest that not only can children learn different causal structures, but they can also use the representations that they have formed in particular ways. Schulz, Bonawitz, and Griffiths (2007) describe children as deriving causal models that "support prediction, intervention, explanation, and counterfactual claims" (p. 1124). It is a key assumption of the influential causal Bayes net approach espoused by Gopnik and colleagues (Gopnik et al., 2004) that children's causal learning should be described in terms of the construction of causal models that can be used to predict hypothetical interventions on variables within a given model (Gopnik & Schulz, 2007; Hagmayer, Sloman, Lagnado, & Waldman, 2007; Sloman & Lagnado, 2005). Extensive reviews of this approach are provided elsewhere (Gopnik et al., 2004; Gopnik & Schulz, 2007), but, put very simply, on the causal Bayes net approach, the causal models that are constructed capture patterns of conditional probability information: i.e., they capture not just the probability of an event occurring, but the probability of events occurring given the occurrence or non-occurrence of other events. For example, a model of a causal chain in which A causes B and B causes C captures the fact that the probability of C occurring given A is not independent of B's occurrence.

A variety of recent studies have explored whether young children are sensitive to such conditional probability information in their causal learning (e.g., Gopnik, Sobel, Schulz, & Glymour, 2001; Kushnir & Gopnik, 2007; Schulz & Gopnik, 2004; Sobel & Kirkham, 2006). As Hagmayer et al. (2007) have demonstrated, because causal Bayes nets summarize conditional probability information, they can be used flexibly to generate predictions about what should happen to a variable in a system if another one is manipulated (given additional assumptions about the nature of causal models). This is held to be one of the major strengths of this approach (see various contributions to Gopnik & Schulz, 2007), and it leads to the prediction that what is learned about the relationships between variables in a causal system should be able to support judgments about the hypothetical or counterfactual effects of intervening on (manipulating the value of) its variables.

This is a claim about what sort of judgments causal representations might be expected to support, but, influenced by the interventionist approach to causation in the philosophical literature (Woodward, 2003, 2007), Schulz, Kushnir, and Gopnik (2007) have made an even stronger claim about what it is to represent a relationship as causal in the first place. Specifically, they argue that "a causal relation is defined...in terms of the real and counterfactual interventions it supports" (p. 69), and that "when children infer that a relationship is causal, they commit to the idea that certain patterns of interventions and outcomes will hold" (p. 70). In other words, what it is to represent a relationship between two variables A and B as causal just is to be committed to certain beliefs about the effects on B of intervening on A (see Woodward, 2007, for discussion).

The aim of the current study is to examine whether judgments about causation and judgments about the effects of interventions are inextricably linked in the way this approach and the causal Bayes net theory might suggest (Schulz, Kushnir, et al., 2007). The empirical prediction we can derive from this theoretical approach is that when a person extracts a particular causal model from their observations of a causal system, they should then be able to make predictions as to how intervening on components of the system will affect the rest of the system (e.g., as modeled by Hagmayer et al., 2007). Importantly, these predictions should be consistent with the causal model they identified in the first place. For example, imagine a participant observed an apparatus such as the one in Figure 1 and has identified the common cause structure illustrated in Figure 1a (A causes both B and C) as the one that shows how the apparatus operates. If they were now asked whether the rectangle (component C) would operate if the circle (component B) was prevented from moving, they should recognize that disabling the circle does not interfere with the effect the star (the A component) has on the rectangle. Likewise, disabling the rectangle would not interfere with the effect the star has on the circle. However, if the participant identified the causal chain structure illustrated in Figure 1b (A causes B and B causes C) as the one that shows how the apparatus operates, then their responses to the two intervention questions should differ. In particular, they should recognize that disabling the circle would also prevent the operation of the rectangle, whereas disabling the rectangle would have no such effect on the operation of the circle. As far as we understand it, a strong interventionist account of causal representation holds that what it is to represent a system as (e.g.,) a causal chain rather than a common cause structure just is to be committed to these differential effects of intervening on these variables.

The interventionist notion of causal representation that Schulz, Kushnir, et al. (2007) subscribe to can be broadly contrasted with one that is influenced by what might be termed a geometrical-mechanical (Woodward, in press) or process (Salmon, 1998) notion of causation (see also Wolff, 2007; Wolff, Barbey, & Hausknetcht, 2010). Loosely speaking, psychological versions of this second approach to causal representation hold that causal relationships are primarily understood not in terms of whether the value of one variable depends on the value of another one, but in terms of there being some type of appropriate physical connection between cause and effect. In developmental psychology, this notion of causation can been seen in Shulz's (1982) claim that children understand causation in terms of force transmission, and in Schlottman's (1999) discussion of the mechanism principle: she claims that when two events are represented as causally related, even preschool children are committed to there being a mechanism by which the cause produces the effect (see also Buchanan & Sobel, in press; Bullock, Gelman, & Baillargeon, 1982). For present purposes, we will use the term "mechanism approach" to causation to contrast with the interventionist approach.

Schulz, Kushnir, et al. (2007) argue against such a mechanism approach, and propose that developmental psychologists should replace it with a notion of causal representation that draws on the interventionist approach. As we have described it, Schulz, Kushnir, et al.'s approach explicitly predicts that if children represent a system as having a certain causal structure, they should be able to make predictions about the effects of intervening on variables in the system. Their interventionist account is a psychological version of a *counterfactual* theory of causation (Harris, German, & Mills, 1996; Hoerl, McCormack, & Beck, in press; Woodward, 2007), in that it directly links causal representation with the ability to think counterfactually about the effects of manipulating variables in a causal system. However, an interesting developmental issue is whether children represent relationships as causal before they can make appropriate counterfactual judgments. Although McCormack, Butterfill, Hoerl, and Burns (2009) found that 4-to-6-year-olds' counterfactual judgments regarding a novel causal property showed the same developmental pattern as children's causal judgments about the same property, other researchers have argued that the ability to think counterfactually may itself emerge relatively late in development (e.g., Beck, Robinson, Carroll, & Apperly, 2007; Rafetseder, Cristi-Vargas, & Perner, 2010). Because a mechanism approach links causal representation to children's grasp of appropriate geometrical-mechanical principles rather than to counterfactual cognition, it leaves open the possibility that children may represent a relationship as causal without being able to make appropriate counterfactual predictions. That is, a mechanism approach at least allows for the possibility of a developmental dissociation between causal judgments and counterfactual judgments. It is important to note, though, that a mechanism approach does not predict that children will be unable to correctly make judgments about the effects of intervening on a causal system. Although, unlike in the interventionist approach, judgments about the effects of intervening on the system are not assumed to be directly given by causal representations themselves, mechanism knowledge might itself be exploited in making counterfactual predictions. On this approach, predictions about the effects of intervening on the system might be assumed to ordinarily depend on what expectations one might have given the nature of the specific mechanisms that connect the system components.

Schulz, Gopnik, et al. (2007, Exp. 2) tested the predictions of the interventionist account by examining children's judgments regarding the effects of intervening on a causal system. In their study, a sample of 16 four-to five-year-old children was told which causal structure was responsible for the operation of a system of two gears. The children were then asked whether each of the gears would operate if the other one was removed. Across the four different causal structures they tested, the children answered the questions correctly more often than was expected by chance. We discuss the methodology of this study in more detail below, but two initial points are worth making about this study and its findings. First, even though performance was above chance, for none of the causal structures did the majority of children (i.e., no more than 8 out of their sample of 16 children) answer correctly. This suggests that children did not find it straightforward to make such judgments. Second, although Schulz, Gopnik, et al. argued that they were assessing whether children could make intervention judgments about a three-variable system, such as those depicted in Figure 1, in fact they only presented children with a two-variable system that had an additional on-off switch. In the study, children were given information that described the dependencies of the two gears (whether the gears did or did not spin by themselves) and they were then were asked whether one gear would work without the other one. Children were only required to make judgments about the relationships between two variables, and thus it is not clear if they could adequately distinguish between the sorts of structures depicted in Figure 1.

----------Insert Figure 1 about here---------

Burns and McCormack (2009) used a system in which children had to reason about dependencies between three variables. In their study, children had to decide which one of the three causal structures shown in Figure 1 they were observing, and then make judgments about how the system would operate if B or C were intervened on. They found that even though 6- to 7-year-olds could effectively discriminate between causal chain and common cause structures, their predictions about the effects on intervening on variables in these systems did not differ appropriately. The current study addresses this issue more systematically, using a wider range of age groups and types of intervention questions (Exps. 1 & 2). Moreover, it also examines whether children can make such judgments under circumstances in which children do not have to initially infer a causal structure (Exp. 3). One important difference between Burns and McCormack's and Schulz, Gopnik, et al.'s (2007) study lies in whether children had to initially decide which causal structure they were presented with, and then make appropriate intervention judgments. In Burns and McCormack's study, children initially made such causal structure judgments, whereas in Schulz et al.'s (2007) study, children were simply told which causal structure obtained. It is difficult to predict under which circumstances children's intervention judgments should be better. If children have had to infer a causal structure, perhaps they may lack confidence in their initial structure judgment and thus guess when making their subsequent intervention judgment. On the other hand, if children have simply been told that a certain structure obtains, rather than reasoning to reach an answer, they may perform more poorly when asked to make intervention judgments because they have not already been asked to reflect on the relationships between the variables. Experiments 1 and 2 explicitly examine the consistency of causal structure and intervention judgments in circumstances in which children have to initially infer a causal structure, whereas Experiment 3 takes a closer look at intervention judgments without requiring children to initially learn the causal structure.

Experiment 1

There is some preliminary evidence to suggest that children can learn whether a structure is a common cause or causal chain if they are shown information about what happens to the variables in a system when other variables are manipulated (Schulz, Gopnik, et al., 2007; Sobel & Sommerville, 2009). However, if children have learned that the structure is of a certain sort by viewing the effects of such manipulations, asking them then to make judgments about the hypothetical or counterfactual effects of intervening on variables may simply test their memories for what they have just observed in the learning phase. Thus, if we want to examine the relationship between children's causal and intervention judgments, it may be better to look at circumstances in which children learn causal structure on the basis of some other type of information. Burns and McCormack (2009) have shown that, at least by 6 to 7 years, children will use simple temporal cues to discriminate between different causal structures (see also Lagnado & Sloman, 2004, 2006). Specifically, they showed that children of this age, and adults, are likely to assume that the events A, B, and C form a causal chain if they consistently see the events occurring in a temporal sequence one after another (A, then B, then C). If they see a different temporal schedule, in which A occurs, and following a delay, B and C occur simultaneously, they are likely to infer that A is a common cause of both B and C. The first two experiments exploit participants' tendency to use simple temporal cues in this way to make causal structure judgments. Thus, in Experiment 1, children in three age groups were initially shown sets of three events with different temporal schedules, and asked to make a causal structure judgment, and then asked to infer how intervention on one component affects the operation of other components.

Method

Participants. We tested 63 children from three school year groups: Group 1 (*N* = 20) were children aged 4-5 years (mean age = 58.25 months; Range = 52 – 63 months) and there were 8 boys and 13 girls. Group 2 (*N* = 21) were children aged 5-6 years (mean age = 70.62 months; Range = 64 – 80 months) and there were 9 boys and 12 girls. Group 3 (*N* = 18) were children aged 6-8 years (mean age = 89 months; Range = 81 – 99 months) and there were 7 boys and 11 girls. Four additional children were excluded from the analysis (because of an equipment malfunction and one boy from Group 1 because he did not complete the task).

Materials. The experimental materials consisted of a 41 cm (long) x 32 cm (wide) x 20 cm (high) wooden box that had three objects inserted on the surface. The box had three predetermined locations for these objects (see Figure 1) which formed an equilateral triangle of sides 24 cm. Any given object could be placed in any of the locations. We used four different colored box lids to signal to the children that they were different systems. The objects were different colored shapes (e.g., circle, square, triangle, crescent) and each of them was only used for one box in a testing session. The objects rotated on the horizontal plane and they were activated by a remote button press that was concealed from the participants. Each object had a small hole drilled in it (see Figure 1), into which a metal rod could be inserted vertically through the hole and into another hole drilled in the box lid, in such a way as to completely prevent the object from moving. This metal rod was fashioned as a "stop sign", which was a miniature version of a red and white road sign with the word "Stop", and when the stop sign was inserted into an object it was perceptually obvious that it could not move. We also installed a dummy on/off switch on the side of the box facing the children, which acted as a setting condition for the initiation of trial sequences. Precise control over the timing of the sequences was ensured by the use of a laptop housed inside the box, of which participants were unaware. Colored photographs of the boxes' surface overlaid with pictures of hands were used in the test phase of the experiment (cf. Schulz, Gopnik, et al., 2007; Sobel & Sommerville, 2009).

Design. The type of causal structure, common cause and causal chain, was a within participants variable and each participant saw two of each, that is, two boxes that had a common cause structure and two that had a causal chain structure (one box ABC and the other ACB). The three different causal structures can be seen in Figure 1. The children were asked to select the picture that showed how each of the boxes worked and were then asked two intervention questions about the B and C components.

Procedure. The children were tested individually and were first introduced to the box and asked to name the color of each of the components. They were then shown three pictures, such as the ones in Figure 1, which were described to them in one of two ways: order of description consistent with causal order or order of description inconsistent with causal order. The order consistent versions all started with a description of the effect object A had on the other objects, e.g.: "In this picture the blue one makes both the black one and the white one go, and the hands show that." (Figure 1a) and "In this picture the blue one makes the white one go and the white one makes the black one go, and the hands show that." (Figure 1b). In the order inconsistent version the descriptions started with object B or C, e.g., "In this picture the white one makes the black one go and the blue one makes the white one go, and the hands show that." (Figure 1b). Order of mention of the components in the picture was varied in this way because the information children were subsequently provided with in order to discriminate between causal structures was temporal in nature. The children were then asked comprehension questions that required them to identify each of the three pictures that had just been shown to them, e.g. "Can you show me the picture where the blue one makes both the white one and the black one go?". When children made errors on the comprehension questions the pictures were described to them again before proceeding with the trial and on the next trial the comprehension questions were asked again.

On completion of the comprehension questions, the children's attention was drawn to the on and off switch at the front of the box and they were asked whether the box was switched on or off (it was always off). They were then told "I am going to switch the box on now and I want you to watch what happens. Remember, you've got to figure out which picture is the right one. Are you ready now?" The children then observed three cycles of the box operating. The temporal schedules were either synchronous (A spins then, following a 0.5 s delay B and C spin simultaneously), or sequential (either A spins then B spins then C spins OR A spins then C spins then B spins, with 0.5 s delay between each event). After observing a box operate three times, the children were asked to identify the picture that shows "how the box really works". Children selected one picture from a set of three (such as those in Figure 1), and their chosen picture was then placed in front of the box and the other two pictures removed from view. The children were then introduced to a new piece belonging to the box, which was the stop sign that resembled a road sign. They were told that it could be used to "stop some parts of the box from going". They were told, for example, that if the stop sign was put into the A object (named by its color) then it couldn't move and this was demonstrated to them. They were then reminded of the picture they had selected as the one that shows how the box works and given one further demonstration of the operating box. They were told “That time the [color A] one moved and the [color B] one and the [color C] one moved.” The stop sign was then inserted into the B (or C object) (counterbalanced) and the experimenter said: “Now imagine I had stopped this one from moving like this. Can you see the [color B] one cannot move anymore? What if that time I had stopped the [color B] from moving would the [color C] one still have moved?”The stop sign was then inserted into the C object and the children were asked a similar counterfactual question about how this would have affected the operation of the B object. Thus, children were asked two intervention questions about each causal system.

The children were then shown the next box and asked to note that this box was different (it was always a different color). Again they were asked to name the colors of the three objects (which were different for each box) and were introduced to the three pictures of the different causal structures, with the pictures showing the particular components appropriate to that box. They were only asked the comprehension questions again if they had failed some or all of them on the first trial. The remaining procedure for the second, third and fourth box were identical. Children saw two boxes with synchronous temporal schedules and two with sequential temporal schedules. The colors of the boxes and the components used for each box were varied across temporal schedules. Children saw the four boxes in one of eight different orders which were pseudo-randomized with the constraint that the first two boxes never had the same temporal schedule.

Results and Discussion

Across the three age groups two thirds of the children required only one comprehension trial. The children in the youngest age group had some difficulties with the comprehension questions and as a result 55% of them required two comprehension trials, whereas this figure was 29% of Group 2 and 20% of Group 3. Thus, the younger children were more likely to need the pictures explained to them more than once, and a chi square analysis on number of children per group requiring either one or two comprehension trials revealed that the differences between the three groups reached statistical significance, $\chi^2(2, N = 59) = 8.56$, $p = .014$.

Subsequent analyses addressed three issues in turn: (i) the nature of children's causal system choices, and how such choices varied by temporal schedule, with the prediction being children should choose a common cause structure for the synchronous temporal schedule and causal chain structures for the sequential schedule; (ii) the responses children gave for the pairs of intervention questions, and whether these varied by temporal schedule in the same manner as causal structure judgments; and (iii) whether children's intervention judgments were consistent or inconsistent with their causal structure judgments.

---------------Insert Table 1 about here-----------------

Causal system choices. The children's picture choices were classified as common cause or causal chain depending on which picture they selected from the set of three pictures (see Table 1 for percentage choices of each as a function of temporal schedule). In addition, in the sequential trials an additional category was used of time-inconsistent: responses were coded as time inconsistent when the children selected ACB for the ABC box and vice versa. Across the three age groups there were no differences in accuracy according to whether the pictures were initially introduced to children in an order consistent or order inconsistent manner, so this variable is not included in any analyses. Children received two trials for each temporal schedule, and Table 2 shows the number of children giving either 0, 1, or 2 responses consistent with each temporal schedule. Chi-square analyses showed that for the youngest group, children's scores did not differ significantly from those predicted by chance χ^2 (2, $N = 20$) = 1.83, $p = 0.40$ for synchronous trials, but did differ from chance for sequential trials χ^2 (2, $N = 20$) = 7.45, $p = 0.02$. For the 5- to 6-year-olds, responses differed from chance for both trial types, $\chi^2(2, N = 21) = 77.6$, $p < .001$ for synchronous and χ^2 (2, N = 21) = 15.5, p < .001 for sequential, and this was also true for the oldest group, χ^2 (2, N = 18) = 111, p < .001 for synchronous trials and χ^2 (2, N = 18) = 36.3, p < .001 for sequential trials. The overall numbers of responses (scores from 0-4) that were consistent with temporal information improved significantly with age, as indicated by a one-way ANOVA, F (2,56) = 8.92, MSE = 1.47, p < .001.

-------------Insert Table 2 about here------------

Intervention questions. Responses to the two intervention questions were categorized according to which causal structure was implied by them. Answering “yes” to both intervention questions was categorized as implying a common cause structure (both B and C would still operate even if the other was disabled); answering “yes” to one intervention question and “no” to the other was categorized as implying a causal chain structure (e.g., B would still operate even if C was disabled, but C would not operate without B). For the sequential schedule, a combination of “yes” and “no” answers was classified as time-inconsistent if it conflicted with the direction of the causal chain indicated by the temporal schedule. Finally, answering “no” to both intervention questions was categorized as coupled (i.e., neither B nor C would operate without each other).

As can be seen from Table 3, across all age groups there was a strong tendency to respond ‘yes’ in response to both the intervention questions regarding whether the not-intervened upon component would move if the adjacent component was prevented from moving. There were no group differences on whether children’s responses to the intervention questions were consistent with the temporal schedules, as tested by a one-way ANOVA, F (2,56) = 1.46, MSE = 2.63, p = .21 on scores 0-4. Across all three age groups, more children gave responses to intervention questions consistent with the temporal schedules for synchronous trials than would be predicted by chance (χ^2 (2, N = 20, 21, 18) > 131, p < .0001 for all groups), but fewer responses to intervention questions consistent with the temporal schedules for sequential trials than would be predicted by chance (χ^2 (2, N = 20, 21, 18) > 9.1, p < .05 for all groups). Thus, although children in all age groups were likely to infer a causal chain structure following the sequential schedule, their responses to the intervention questions did not similarly imply a causal chain.

-------------Insert Table 3 about here-------------

Consistency of responses. Although we have categorized the responses to intervention questions in terms of whether they were consistent with the temporal schedule, this description is perhaps misleading. This is because although the temporal cues could be taken as a cue to causal structure, in fact different causal structures could nevertheless obtain (e.g., the correct causal structure following sequential schedule, in which A, B, and C occurred in a sequence, could nevertheless be a common cause one, in which A independently causes both B and C, but the delay between A and the occurrence of B may be shorter than the delay between A and the occurrence of C). Thus, it would not be incorrect for children to infer a common cause structure following a sequential temporal schedule, nor to then give answers to the intervention questions that would be categorized as common cause. The consistency of children's responses to intervention questions that is important is actually their consistency with children's initial causal structure choices, not their consistency with the temporal schedules. For example, if children have selected a common cause structure, will they give the appropriate answers to the intervention questions ("yes" and "yes"), or if they have selected a causal chain structure, will they appropriately answer "yes" and "no" to the intervention questions? Table 4 shows the percentage of responses to intervention questions that were consistent with children's initial causal structure choices for each age group, according to causal model choice. It can be seen from the table that children tend to answer the intervention questions appropriately if they have selected a common cause model. However, this would seem to reflect an overall bias to respond "yes" to the intervention questions, rather than a genuine appreciation of the implications of their causal structure choice. When children indicated that they thought the structure was a causal chain (either

ABC or ACB), they rarely answered the intervention questions in a manner that was consistent with their chosen causal structure.

--------------Insert Table 4 about here-----------------

Table 5 gives the distributions of scores (0-2) for responses that were consistent with children's initial causal structure choices for each of the temporal schedules. These distributions do not differ significantly from chance for 4- to 5-year-olds. For the 5- to 6-year-olds, the distributions differ from chance for the synchronous schedule, χ^2 *(2, N = 21)* = 48, *p* < .001, but not for the sequential schedule, χ^2 (2, *N* = 21) = 1.81, *p* = .40. For the oldest group, the distributions differed from chance for both temporal schedules, χ^2 (2, *N* = 18) = 134, *p* < .001, χ^2 (2, *N* = 18) = 24.5, *p* < .001 for the synchronous and sequential schedules respectively.

----------Insert Table 5 about here----------

To summarize, there was a clear developmental trend in extent to which children used the simple temporal cues to extract causal structure. The 4-5 year olds tended to use the sequential temporal structure to infer the appropriate causal chain more often than would be expected by chance, but did not infer a common cause causal structure when the temporal schedule was synchronous. As we have pointed out, it is not necessarily incorrect to infer a causal chain for a synchronous schedule. However, older children did indeed use the temporal information in the synchronous schedule to infer a common cause structure, in a similar way to adults (Burns & McCormack, 2009). We would argue on the basis of these, and other findings, that temporal cues provide a highly salient basis on which to make causal structure judgments. Indeed, in some circumstances such cues appear to be more salient than covariation information (Lagnado & Sloman, 2004, 2006; Frosch, McCormack, & Lagnado, 2011), although, in the absence of temporal information, covariation information is exploited in making causal structure judgments (Frosch et al., 2011; Schulz, Gopnik, et al., 2007; Steyvers et al., 2003). However, the more important result from this experiment was that children in all age groups did not reliably answer the intervention questions in a manner that was consistent with the causal structure that they had initially selected. Although older children did so more often that would be expected by chance, this was because of a tendency to respond “yes” to both intervention questions, a response pattern consistent with the common cause model. Correct responses when a causal chain model was selected were rare.

Experiment 2

In Experiment 2 we tested the oldest age group in the previous experiment in order to further examine their ability to answer questions about intervention. The 6- to 7-year-olds in the previous experiment reliably used the temporal cues to identify causal structure and so we focused on modifying the second part of the task in an attempt to improve children’s performance. The first change was to the way the intervention questions were phrased. In Experiment 1, children were asked counterfactual questions about what would have happened if one of the components had not been able to operate. We used counterfactual questions in the light of the long-standing debate about the relationship between causal and counterfactual judgments (Hoerl et al., in press; McCormack, et al., 2009; McCormack, Frosch, & Burns, in press). Moreover, McCormack et al. (2009) have demonstrated that children of a similar age can make sense of counterfactual questions about scenarios involving novel causal powers and unfamiliar objects. Nevertheless, in Experiment 2, instead of asking counterfactual questions about what would have happened, we asked future hypothetical questions, e.g. ‘If I switch the box on now will the [other shape] go?’. This modification was made on the basis of research findings that suggest that children may find future hypothetical questions (“what will happen”) easier than counterfactual questions (“what would have happened”) (Perner, Sprung, & Steinkogler, 2004; Riggs, Peterson, Robinson, & Mitchell, 1998).

The second change related to the intervention questions. The intervention questions posed in Experiment 1 required two inferential steps and can be described as prevent-then-generate interventions. The first step requires the participant to infer the consequences of preventing one event (B or C) from occurring and the second step requires the participant to imagine the entire event sequence when event A is initiated. Arguably, such interventions are potentially more difficult to think about than an intervention which requires only one step, such as simple generative interventions (e.g., what would happen if B was manipulated?). Therefore, in addition to asking a prevent-then-generate intervention question, we also asked a generative intervention questions, such as: "If I move the white one round like this will the black one go?". Experiment 2 also differed from Experiment 1 in that interventions to disable the components were made by the experimenter actually manually preventing the component from operating, rather than using the stop sign. This was to make this intervention more similar to the generative intervention, in which the experimenter was to intervene by making a component move, rather than by disabling a component. The final difference was that children now only saw one trial involving a single temporal schedule, thus shortening the task and ensuring that children could not be confused by the other causal systems that they had observed.

Method

Participants. Sixty 6-to-7-year-olds participated in the study (*M* = 86 months; *Range* = 79 – 92 months). There were 32 girls and 28 boys. Children were randomly assigned to one of the two experimental conditions.

Apparatus. The apparatus used was similar, though not identical, to the one used in Experiment 1. The box measured 33 cm x 45 cm x 15cm and the three objects formed an equilateral triangle of sides 32 cm. Only three objects were used in this experiment because participants were shown a single causal system; these objects were a blue ball sitting on top of a bent spindle which rotated in the horizontal plane along an elliptical pathway, a yellow square (6 cm x 6 cm) and a red cylindrical bar (10.5 cm long and 2.5 cm in diameter). The location that each component occupied was varied across participants. Similar pictures of the box to those used in Experiment 1 were used at test to elicit children's causal structure choices.

Design. The design was similar to that of Experiment 1, with the on/off switch on the front of the box used as a setting condition for the autonomous activation of the event sequences. The temporal schedules that event sequences followed were also the same as that used in Experiment 1. However, in this experiment causal structure was a between subjects factor. Accordingly, all participants were only shown either a sequential sequence or a synchronous sequence.

Procedure. The procedure closely followed that of Experiment 1. However, after the initial introduction to the box children observed four instead of three demonstrations of the box's operation. Each one began shortly after the box was switched "on" by the experimenter. After the occurrence of all three events on each demonstration the experimenter switched the box "off" again. The experimenter then presented children with the same three causal structure pictures as used in Experiment 1. The causal structure pictures were placed in a random order between the child and the box. Each picture was described in turn by the experimenter in the same manner as in Experiment 1. Children then observed two further demonstrations, after which they were asked to select the causal structure picture that they thought showed how the box actually worked.

The selected picture remained in full view of the child in front of the box while the other two pictures were removed out of sight. Children were then asked two prevent-then-generate intervention questions phrased as future hypotheticals. The experimenter prevented either event B or C in turn from occurring by holding it firmly and saying to children "I am going to stop this one from going *(intervention was then performed)*. So now this one cannot go at all. It really cannot go", and then asked children "If I switch the box on now will the [e.g., blue] one go?" Finally, having answered these two prevent-then-generate intervention questions, children were asked two generative intervention questions. These were of the form 'If I make B go will C go?' and 'If I make C go will B go?' The questions were phrased to children as follows: "If I move the red one around like this, will the yellow one go?” The experimenter mimicked moving the device with his hand, above the device itself. As with the preventative intervention questions the order in which the two questions were asked was counterbalanced across participants, however, all children were asked the prevent-then-generate intervention questions before the generative intervention questions.

Results and Discussion

Table 6 shows the percentage of children who gave each type of response to the causal model question, prevent-then-generate intervention questions and the generative intervention questions, using the categorizations introduced in the previous experiment.

--------------Insert Table 6 about here--------------

Causal model Choice. Analysis revealed that there was a significant difference in the distribution of common cause versus causal chain causal model picture choice across the two conditions, χ^2 (1, N = 60) = 17.68, p < 0.001. The majority of children in the synchronous condition selected the common cause model while the majority of children in the sequential condition selected the causal chain model. Only 7% of participants in the sequential condition gave time-inconsistent responses in the sequential condition (i.e., selected the causal chain model inconsistent with the temporal order of events).

Prevent-then-generate Intervention Questions. Although common cause responses were the modal response in both conditions, there were a large number of coupled responses in the synchronous condition and time-inconsistent type responses in the sequential condition. Thus, unlike in Experiment 1, common cause responses (answering “yes” to both questions), although frequent, did not form the majority of responses. In the synchronous condition, more children gave common cause responses than would be expected by chance (binomial test, p < .001). However, in the sequential condition, the number of children who gave appropriate causal chain responses to the intervention questions did not differ from that expected by chance (binomial test, $p > .10$).

Generative Intervention Questions. Only a minority of children gave common cause or causal chain type responses across the two conditions. Coupled responses were the modal and majority response in the synchronous condition (i.e., judging that moving B would make C move, and also that moving C would make B move). The number of children who gave common cause responses in the synchronous condition did not differ from that expected by chance, nor did the number of children who gave causal chain responses in the sequential condition (binomial test, $p > .16$ in both instances).

Consistency of Response. As we have said in the discussion of Experiment 1, what is critical is the extent to which children give responses to the intervention questions that are consistent with their causal structure choices, rather than consistent with the temporal schedules. The percentage of children who answered the prevent-then-generate intervention questions consistent with their choice of causal model picture was just 38%. Even though more children across both conditions gave common cause responses to the prevent-then-generate intervention questions than responses of any other type, they were not significantly more likely to be consistent in the synchronous condition than in the sequential condition, although the chi square approached significance, χ^2 (1, N = 60) = 3.46, p = 0.06. The consistency of response between causal model picture choice and the generative intervention questions was just 30%. This is unsurprising as many children in both conditions said ‘yes’ to both generative intervention questions, which was not consistent with any of the causal model picture choices that they had made.

Thus, as in Experiment 1, the majority of 6- to 7-year-old children were guided in their causal structure choices by the temporal schedules across the two conditions. Their performance on the intervention questions was still poor, and likely to be inconsistent with their causal structure choices, despite the fact that they were asked future hypothetical rather than counterfactual questions, and despite simplifying the task so that children saw only one box and made only one set of intervention judgments. Children's responses on the prevent-then-generate intervention questions were not dominated by common cause responses (i.e., answering yes to both questions), as they were in Experiment 1. In this experiment, children gave more responses of other types to these questions than common cause responses. This suggests that their poor performance in Experiment 1 on intervention questions was not simply due to a "yes" bias. Nevertheless, they were still no more likely than chance to give causal chain responses to the intervention questions in the sequential condition. Moreover, performance on the generative intervention questions, which we thought might be simpler, was no better than in the prevent-then-generate intervention questions. In fact, children were likely to judge that moving either B or C would make the other component move regardless of whether they had selected a common cause or causal chain structure.

Experiment 3

The findings from the first two experiments suggest that even 6- to 7-year-old children find it difficult to make intervention judgments that are consistent with their causal judgments. However, this conclusion stands in contrast to that drawn by Schulz, Gopnik, et al. (2007) on the basis of the findings of their second experiment; those authors conclude that even 4-year-olds can make appropriate intervention judgments. There are a number of differences between the methodology used by Schulz, Gopnik, et al. (2007) and the one used in our studies reported so far, such as the different apparatus, and we return to these differences in more detail in the General Discussion. The main difference we focus on here is that the children in our experiments had to identify a causal structure themselves and were then asked to make inferences from the structure they had identified, whereas Schulz, Gopnik, et al. told their participants what the structure was. The children in our experiments inferred a causal structure based on temporal cues. One might argue that their causal structure choices simply reflect their perceptions of the temporal relations between the events rather than a genuine understanding of the underlying causal structure (see McCormack et al., in press, for discussion; although Burns & McCormack, 2009, found that the intervention judgments of adults were consistent with their causal structure judgments in a similar procedure). If this is the case, then it would not be surprising if their intervention judgments were not consistent with their causal structure judgments. In order to rule out this explanation of children's relatively poor performance on intervention questions, in Experiment 3, we followed Schulz, Gopnik, et al. in telling children what the causal structure was before asking them the intervention questions rather than asking them to make inferences. We also followed Schulz, Gopnik, et al. in actually demonstrating the effects of an intervention to children during the introduction to the test trials, which we had not done in the previous two experiments, and we worded the intervention questions in the same way as they had.

Method

Participants. We tested 68 children from two age groups: Group 1 were 37 4- to 5-year-olds, 20 boys and 17 girls (mean age = 64.59 months, range = 58-69 months) and Group 2 were 31 6- to 7-year-olds, eight boys and 23 girls (mean age = 86.16 months, range = 79-91 months).

Materials. These were identical to those used in Experiment 1, including the stop sign used to disable objects.

Design. Each child was presented with a training trial followed by three trials where they were presented with three different causal structures, one common cause, one causal chain ABC, and one causal chain ACB, as depicted in Figure 1. Across participants, the three causal structures were presented in six different orders.

Procedure. In keeping with the procedure used by Schulz et al. (2007), the children were introduced to the various components and functionality of the box during a training phase. Using only two objects, we demonstrated that the objects could be removed from the box and children were told that "some objects spin by themselves and some need others to make them move".

Children were initially asked to name the colors of the objects to check that they knew the appropriate color words. We then introduced the stop sign and demonstrated that it would stop a shape from spinning as well as preventing other shapes from spinning. Each child observed the following sequence of events involving the two training objects, with shapes referred to by their colors:

1. Both shapes were initially removed from the box. One shape was added to the box, and that shape moved once the box was switched on, demonstrating that “some shapes move on their own”.
2. The second shape was added, and then both shapes moved once the box was switched on (the first shape spun, followed by the rotation of the second shape), and children were told that the second shape was moved by the first one. This demonstrated that “some shapes are moved by other shapes”.
3. The stop sign was inserted into first shape, and then neither shape rotated when the box was switched on, to demonstrate that the stop sign would prevent one shape from moving, and thus prevent its effect from also occurring. Children were told that the second shape did not move because it needed the first shape (which was disabled) to make it move: “See the [color] one didn’t go because it needs the [other color] one to make it go”.

The children were then introduced to a new box with a new set of three components. As in Experiment 1, the children were first introduced to the causal model pictures that illustrated the different ways in which the box might work. They were also asked the comprehension questions used in previous experiments to verify they had understood the differences between the three pictures. The descriptions of what the pictures depicted and the comprehension questions were repeated if necessary, until the child answered all comprehension questions correctly. The box was switched on and they watched it operate three times. For the common cause structure, the temporal schedule of the objects’ movement was a simultaneous one, and for the causal chain structures the temporal schedule was the sequential schedule appropriate to the direction of the causal chain. The experimenter then showed them the particular picture that illustrated how the first box worked and described the picture to them. The other two pictures were removed from view. The children were then reminded of the stop sign's function: "It can be used to stop some parts of the box from going. For example, if I put it into the [color A] one then it won't move." They were then reminded of the picture that showed how the box works (the picture was re-described to them, e.g., "The blue one makes the white one spin, and the white one makes the black one spin"), and given one further demonstration of the box's operation. The stop sign was then inserted into the B or C object and the other object was lifted above its peg. The experimenter then said: "If I put the [color C] one down right now and turn on the switch, will the [color C] spin or stay still?" The same procedure was then followed for the other object. Following this trial, a new box and set of objects was introduced to children, along with the appropriate picture of causal structure. The order in which children were shown each causal structure was varied across participants.

Results

The children in the youngest age group had some difficulties with the comprehension questions and as a result 49% of them required more than one comprehension trial (22% required 2 trials, 19% required 3 trials and 8% required 4 trials). The 6- to 7-year-old children had few difficulties with the comprehension questions with only 26% requiring more than one trial. A chi square analysis on number of children per group requiring more than one comprehension trials revealed that the difference between the two groups was marginally statistically significant; χ^2 (1, N = 68) = 3.73, p = .054.

---------------Insert Table 7 about here-------------

Table 7 provides a breakdown for how the responses could be classified in relation to the different causal structures that they imply. The percentage of children in the younger age group that answered both of the causal structure intervention questions correctly was 8% on the synchronous trials, 16% on the ABC trials and 27% on the ACB trial, and for none of these trials did children give the appropriate response to the intervention questions more often than would be predicted by chance (binomial test, $p > .10$). The proportion of children in the older age group that answered each of the causal structure intervention questions correctly was 32% on the synchronous trials, 43% on the ABC trials and 52% on the ACB trials. The older children did not give the appropriate answer more often than would predicted by chance in the common cause trial, but, unlike in the previous two experiments they did do so in each of the causal chain trials (binomial test, $p < .02$ for both chains). However, we note that even though as a group children performed above chance on the causal chain trials, this group answered the intervention questions about causal chains correctly only 48% of time. Thus, overall level of performance was still low. Table 8 shows the distribution of scores (0-3) across the three trial types for each age group. This distribution differs from chance for the 4-5-year-olds, but because children were performing worse than would be expected by chance, χ^2 (3, N = 37) = 11.9, p = .01. This was due to the predominance of coupled responses. However, the older children got more trials correct than would be predicted by chance, χ^2 (3, N = 30) = 24.4, p < .001.

-------------Insert Table 8 about here-----------

It is worth pointing out that children in both groups were more likely than chance to give coupled responses (i.e., to answer “no” to both intervention questions) when told that the causal structure was a common cause structure. It might be argued that this response is not necessarily incorrect, in that although children have explicitly been told that A causes B and also that A causes C, and shown the appropriate picture, this information is compatible with a belief that B and C nevertheless will not work without each other. However, it is not possible to make a similar argument in the case of the causal chain causal structures. For those structures, answering that B will not work with C and C will not work without B could not be construed as a correct response.

Not only were children told that (e.g.,) A makes B go, and B makes C go, but they also saw the events in a temporal sequence ABC. Judging that B would not operate without C would involve ignoring what they had been told about the causes of B and C and also the implications of the temporal order in which events occurred (because C did not happen until after B had already occurred). We also note that the most common response in this experiment was to answer “no” to intervention questions, which stands in contrast to the modal response in Experiment 1, in which “yes” responses were common. Thus, the findings of this study cannot be explained in terms of bias to respond “yes”.

General Discussion

The most striking finding from the studies reported here is that across all age groups of children tested, intervention judgments were not reliably consistent with causal ones. This finding contrasts with those of Schulz, Gopnik, et al. (2007) and is not predicted by a strong interventionist account of causal representation. For example, across the three experiments, for the oldest age-group tested, even though children either judged (Exps. 1 & 2) or were told (Exp. 3) that the causal structure was (e.g.,) a causal chain ABC, only around 30-50% of the time did they appropriately judge that C would not operate without B but that B could operate without C. This was despite the fact that they always saw the events occur in a temporal sequence in which C occurred after B. Although in Experiment 1, poor performance on intervention questions appeared to be due to a general tendency to give positive responses to questions about the effects of intervening on a component, such a response bias did not dominate in Experiments 2 and 3 and performance remained at similar levels. The performance of the older children was significantly better than chance on some causal structures in the third experiment, but nevertheless children gave responses to intervention questions that were inconsistent with causal structure more often than consistent responses.

It could be argued that the basis of the causal judgments in Experiments 1 and 2 was such that they did not support intervention judgments easily. Causal structure judgments were based solely on simple temporal cues, and perhaps this led to them having a kind of fragility that meant that they did not support intervention judgments. It may be that intervention judgments are easier to make if causal structure judgments have been arrived at by observation of patterns of contingency information. For example, if you have never seen C happen in the absence of B, then you might be more confident about making a judgment about the effect of intervening on B. Indeed, it would be interesting to examine whether the consistency of causal and intervention judgments depends on the basis on which causal judgments are made. In Experiment 3, we followed Schulz, Gopnik, et al. (2007) in actually telling children what the causal structures were. This, and other modifications we made to the task in Experiment 3, seemed to improve performance for the older children but not the younger children. However, a substantial number of the older children still found the task difficult and thus answered incorrectly.

How can our results be reconciled with Schulz, Gopnik, et al.'s (2007) conclusion that "children can use knowledge of causal structure to predict the pattern of evidence that will result from interventions" (p. 328)? As we have already mentioned, whilst the 16 children who participated in Schulz, Gopnik, et al.'s study performed above chance, not more than 50% of their children answered correctly for any given causal structure. This level of performance is comparable to the performance of the older children in our third experiment. As pointed out in the introduction, a key difference between our study and that of Schulz, Gopnik, et al. (2007) was that we used diagrams of a three variable system that required children to grasp the nature of the relationships between all three variables, whereas Schulz, Gopnik, et al. (2007) asked children to focus on the relationships between two variables in a system (i.e., they only depicted B and C in their diagrams, and children were only required to reason about the causal relationship between these two variables). While we accept that explicitly including all three variables may have made our task more complex for children, we would argue that it is necessary to do so if we want to examine whether children can represent and reason about the sorts of causal structures in question – common cause and causal chain structures – which by definition must include at least three variables.

Our study and that of Schulz, Gopnik, et al. (2007) also differed in two ways we have not considered thus far. In the Schulz, Gopnik, et al. study the objects were gears that came into physical contact with one another because they interlocked when they spun. The objects in our system were spatially separated. We used spatially separated objects because it seemed to us that, given the nature of gears, it was not actually physically possible for two objects in Schulz, Gopnik, et al.'s (2007) study to "spin by themselves", although their experimental procedure involved children being told that this was the case. That is, if gears are interlocking and one gear moves, the other will have to move as well. By using spatial separated objects, we ensured children could at least make sense of the description of some objects as moving by themselves, so if anything this procedure difference might be expected to make our task easier. A second important difference between the two studies was that Schulz and her colleagues intervened on a component by removing it from the causal system entirely, whereas we intervened on a component by preventing it from moving (by inserting the stop sign in Exps. 1 and 3). Formally, completely removing a component can actually change the causal system to become a new one, whereas disabling a component involves making an intervention on a variable within the existing system.[1] We do not know if children's performance on intervention questions might depend on the specific type of intervention that is performed. However, the results of Experiment 2, when children were asked not only questions about how the system will operate if B is disabled, but also generative intervention questions – questions about whether making B or C happen will make the other component move – suggest that children do not necessarily find some sorts of intervention questions easier than others.

We know from our own studies with adults that in this type of task the majority of adults' intervention judgments are consistent with their causal ones (Burns & McCormack, 2009). Moreover, we did observe developmental changes in children's intervention judgments, with older children being more likely to give judgments consistent with causal structure than younger children. However, while performance of the oldest children was above chance in some cases, it is still not clear evidence that the children's causal representations easily support inferences about intervention. Instead, we would argue that we have observed a development towards these skills and that further research is required to fully understand how these skills develop. As it stands, the evidence is not sufficient to support the strong interventionist claim that "when children infer that a relationship is causal, they commit to the idea that certain patterns of interventions and outcomes will hold" (Schulz, Kushnir, et al., 2007, p. 70). Instead, it appears that the ability to recognize that a relationship is causal and an ability to make inferences about the effect of interventions on this relationship are two skills that are not straightforwardly linked. It is possible that, when faced with a novel mechanical system, such as the one used in the current study, children may find it difficult to give coherent answers about the effects of intervening on components of the system without knowing anything about the underlying mechanisms that connect the components. Such a suggestion would be consistent with a mechanism approach to causal representation.

Importantly, though, Woodward (in press) has argued that what we have termed a mechanism approach to causation and an interventionist approach to causation should be considered to capture quite different aspects of our causal cognition, rather than viewed as competing theories of the nature of causal representation. He argues that in mature adults, these two ways of thinking about causation co-exist and are typically well-integrated. Thus, adults can move from information about mechanism to inferences about interventions and dependency relationships and vice versa. It may be that adults' responses about interventions are consistent with their causal structure responses precisely because they explicitly recognize the need for such consistency. However, Woodward suggests that in some populations, such as young children or animals, these two types of information may yet to be properly integrated. Thus, we might expect to see some sorts of dissociations between aspects of their behavior depending on how causal cognition is assessed. Thus, rather than interpreting the current data as suggesting that an interventionist approach is incorrect, it may be more useful to consider them as providing grounds for considering how the interventionist and mechanism components of causal cognition may change developmentally, and how these two types of ways of thinking about causation become properly integrated. Clearly, this is an area in which additional research is warranted, due to its important implications for theories of causal cognition.

References

Beck, S. R., Robinson, E. J., Carroll, D. J., & Apperly, I. A. (2006). Children's thinking about counterfactuals and future hypotheticals as possibilities. *Child Development, 77,* 413-426.

Buchanan, D. W., & Sobel, D. (in press). Mechanism-based causal reasoning in young children. *Child Development.*

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York: Academic Press.

Burns, P. & McCormack, T. (2009). Temporal information and children's and adults' causal inferences, *Thinking & Reasoning, 15*, 167-196.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Frosch, C., McCormack, T., & Lagnado, D. A. (2011). *Temporal influences on causal structure judgments.* Manuscript submitted for publication.

Gopnik, A., Glymour, C., Sobel, D.M., Schulz, D.E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 1-31.

Gopnik, A. & Schulz, L. (2007). (Eds.), *Causal learning: Psychology, philosophy, and computation*. Oxford, England: Oxford University Press.

Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology*, *37*, 620-629.

Hagmayer, Y., Sloman, S., Lagnado, D. A., & Waldman, M. R. (2007). Causal reasoning through intervention. In A. Gopnik & L. E. Schulz (eds.), *Causal learning: psychology, philosophy, and computation* (pp. 86-100). Oxford, England: Oxford University press.

Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition, 61,* 233-259.

Hoerl, C., McCormack, T., & Beck, S. R. (in press). (Eds.), Understanding counterfactuals, understanding causation: Issues in philosophy and psychology. Oxford: Oxford University Press.

Kushnir, T. & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology, 43*, 186-196.

Kushnir, T., Gopnik, A., Lucas, C., & Schulz, L. (2010). Inferring hidden causal structure. *Cognitive Science, 34,* 148-160.

Lagnado, D. A. & Sloman, S. A. (2004). The advantages of timely intervention. *Journal of Experimental Psychology: Language, Memory, and Cognition*, *30*, 856-876.

Lagnado, D. A. & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Language, Memory, and Cognition*, *32*, 451-460.

McCormack, T., Butterfill, S., Hoerl, C., & Burns, P. (2009). Cue competition effects and young children’s causal and counterfactual reasoning. *Developmental Psychology, 45,* 1563-1575.

McCormack, T., Frosch, C. A., & Burns, P. (in press). The relationship between children’s causal and counterfactual judgments. In C. Hoerl, T. McCormack, & S. R. Beck (Eds.), *Understanding counterfactuals, understanding causation: Issues in Philosophy and Psychology*. Oxford, England: Oxford University Press.

Perner, J., Sprung, M., & Steinkogler, B. (2004). Counterfactual conditionals and false belief: a developmental dissociation. *Cognitive Development*, *19*, 179-201.

Rafetseder, E., Cristi-Vargas, R., & Perner, J. (2010). Counterfactual reasoning: Developing a sense of “nearest possible world”. *Child Development, 81,* 376-389.

Riggs, K. J., Peterson, D. M., Robinson, E. J., & Mitchell, P. (1998). Are errors in false belief tasks symptomatic of a broader difficulty with counterfactuality? *Cognitive Development, 13*, 73-90.

Schlottmann, A. (1999). Seeing how it happened and knowing how it works: How children understand the relationship between perceptual causality and underlying mechanism. *Developmental Psychology, 35,* 303-317.

Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? Naïve theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology, 43,* 1124-1139.

Schulz, L. & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40,* 162-176.

Schulz, L., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science, 10*(3) 322-332.

Schulz, L., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Intervention and causal inference. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 67-85). Oxford: Oxford University Press.

Shanks, D. R. (2004). Judging covariation and causation. In Koehler, D.J., and Harvey, N. (Eds.) *Handbook of judgment and decision making*. (pp. 220-239) Oxford, England: Blackwell.

Shultz, T. R. (1982). Rules for causal attribution. *Monographs of the Society for Research in Child Development, 47*(1, Serial No. 194).

Sloman, S. A. & Lagnado, D. A. (2005). Do we "do"? *Cognitive Science, 29,* 5-39.

Sobel, D. M. & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42,* 1103-1115.

Sobel, D. M. & Kushnir, T. (2006). The importance of decision demands in causal learning from intervention. *Memory & Cognition, 34*, 411-419.

Sobel, D. M. & Sommerville, J. A. (2009). Rationales in children's causal learning from others' actions. *Cognitive Development, 24,* 70-79.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science, 27,* 453-489.

Woodward, J. (2003). *Making things happen: A theory of causal explanation.* Oxford, England: Oxford University Press.

Woodward, J. (2007). Interventionist theories of causation in psychological perspective. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 19-36). Oxford, England: Oxford University Press.

Woodward, J. (in press). A philosopher looks at tool use and causal understanding. In T. McCormack, C. Hoerl, & S. Butterfill (Eds.), *Tool use and causal cognition.* Oxford: Oxford University Press.

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General,* 136, 82-111

Wolff, P., Barbey, A. K., & Hausknecht, M. (2010). For want of a nail: How absences cause events. *Journal of Experimental Psychology: General, 139,* 191-221.

Author Note

This project was funded by an ESRC grant (Grant No. RES-062-23-1799) awarded to Teresa McCormack and David Lagnado.

Footnote

[1] This is clear if we think of an example: imagine a system in which there are three billiard balls in a line, A, B, and C, such that hitting A leads to it striking B which then strikes C. We could intervene on this system either by removing B entirely, or by disabling B (e.g., gluing it to the table). Removing B entirely essentially changes the causal system in play, whereas we could predict the effects of disabling B by assuming that the system is a causal chain.

Table 1. *Percentage data for type of causal structure picture identified according to temporal schedule and group in Experiment 1. Shaded cells are the responses predicted if temporal cues are used.*

		Picture choice			
%		Common cause	Causal chain	Time inconsistent	Don't know
4- 5 yrs	Synchronous	37.5	60	0	2.5
N = 20	Sequential	30	45	25	
5-6 yrs	Synchronous	81	19		
N = 21	Sequential	36	55	10	
6-8 yrs	Synchronous	94	6		
N = 18	Sequential	19.5	69.5	11	

Table 2. *Number of Causal Structure Choices (out of 2) Consistent With Temporal Schedules For Each Age Group and Schedule in Experiment 1.*

		Number of responses consistent with temporal schedule		
		0	1	2
4- to 5-year-olds	Synchronous	9	7	4
	Sequential	8	6	6
5- to 6-year-olds	Synchronous	2	4	15
	Sequential	6	7	8
6- to 8-year-olds	Synchronous	0	2	16
	Sequential	3	5	10

Table 3. *Percentage Data for Causal Structure Implied by Response to the Two Intervention Questions According to Temporal Schedule and Group in Exp. 1. Shaded cells are the Responses Implied by the Temporal Schedule.*

		Causal structure implied by responses				
%		Common cause	Causal chain	Time inconsistent	Coupled	Don't know
4-5 yrs	Synchronous	85	3	-	8	5
	Sequential	85	5	7		3
5-6 yrs	Synchronous	71	2	-	26	
	Sequential	67	5	26		2
6-8 yrs	Synchronous	81	8	-	11	
	Sequential	61	28	3	8	

Table 4. *The Percentage of Times Children Gave Responses to the Intervention Questions That Were Consistent With Their Causal Structure Choice in Experiment 1.*

	Common cause	Causal chain ABC	Causal chain ACB
4- to 5-year-olds	74%	5%	0%
5- to 6-year-olds	61%	0%	12%
6- to 8-year-olds	81%	22%	26%

Table 5. *Number of Responses to Intervention Questions (out of 2) Consistent With Causal Structure Choices For Each Age Group and Schedule for Experiment 1.*

		Number of responses consistent with causal structure choices		
		0	1	2
4- to 5-year-olds	Synchronous	11	6	3
	Sequential	12	7	1
5- to 6-year-olds	Synchronous	7	5	9
	Sequential	14	5	2
6- to 8-year-olds	Synchronous	3	2	13
	Sequential	10	2	6

Table 6. *Percentage of Children Who Gave Each Response Type to Each Question in Experiment 2*

		Synchronous	Sequential
Causal Model	Common Cause	73%	20%
	Causal Chain	27%	73%
	Time-inconsistent	-	7%
Prevent-then-generate Intervention Questions	Common Cause	50%	43%
	Causal Chain	20%	17%
	Coupled	30%	30%
	Time-inconsistent	-	10%
Generative Intervention Questions	Common Cause	27%	33%
	Causal Chain	13%	23%
	Coupled	60%	43%
	Time Inconsistent	-	0

Table 7. *Percentage of Responses to Intervention Questions Implying the Different Causal Structures in Experiment 3.*

%	Trial	Common cause	Causal chain: consistent	Causal chain: inconsistent	Coupled
4-5 yrs *N* = 37	Common cause	8	16	-	76
	Causal chain ABC	3	16	13.5	67.5
	Causal Chain ACB	0	27	13.5	59.5
6-7 yrs *N* = 31	Common cause	32	10	-	58
	Causal chain ABC	10	43	7	40
	Causal chain ACB	16	52	0	32

Table 8. *Distribution of Numbers of Correct Trials (0-3) in Experiment 3.*

	Total number of correct trials			
	0	1	2	3
4- to 5-year-olds	25	6	5	1
6- to 7-year-olds	8	9	10	3

Figure Caption

Figure 1. The three different causal structures displayed by the boxes: a) common cause – A causes both B and C, b) causal chain 1 – A causes B and B causes C, C) causal chain 2 – A causes C and C causes B.

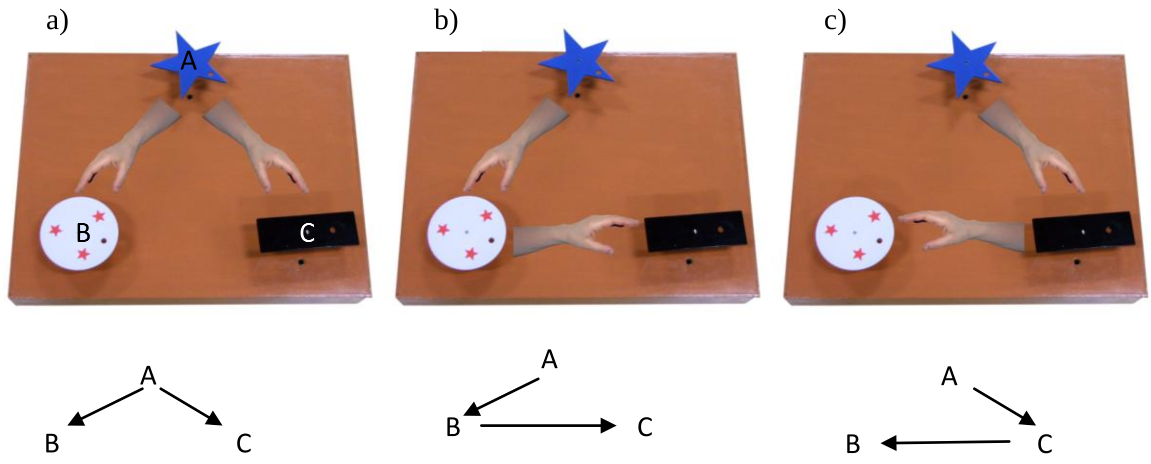
a)
b)
c)
A
B
C
A
B
C
A
B
C
A
B
C