



Children Use Temporal Cues to Learn Causal Directionality

Benjamin M. Rottman,^a Jonathan F. Kominsky,^b Frank C. Keil^b

^a*Section of Hospital Medicine, Department of Medicine, University of Chicago*

^b*Department of Psychology, Yale University*

Abstract

The ability to learn the direction of causal relations is critical for understanding and acting in the world. We investigated how children learn causal directionality in situations in which the states of variables are temporally dependent (i.e., autocorrelated). In Experiment 1, children learned about causal direction by comparing the states of one variable before versus after an intervention on another variable. In Experiment 2, children reliably inferred causal directionality merely from observing how two variables change over time; they interpreted Y changing without a change in X as evidence that Y does not influence X . Both of these strategies make sense if one believes the variables to be temporally dependent. We discuss the implications of these results for interpreting previous findings. More broadly, given that many real-world environments are characterized by temporal dependency, these results suggest strategies that children may use to learn the causal structure of their environments.

Keywords: Causal structure; Causal learning; Causal direction; Time; Observation; Intervention

1. Introduction

The ability to identify the direction of causal relations is critical for explaining events and taking actions. For example, correlations between “brain” diseases (e.g., Parkinson’s disease and autism) and intestinal disease have recently been identified, with some evidence suggesting that the gut pathology might cause the brain pathology (Forsyth et al., 2011; White, 2003). Knowing the true direction of these causal relations not only leads to a more accurate understanding of the diseases but also facilitates the development of

Correspondence should be sent to Benjamin M. Rottman, Department of Psychology, LRDC 726, 3939 O’Hara St., Pittsburgh, PA, 15260. E-mail: rottman@pitt.edu

treatments. Similarly, in more mundane daily events experienced by children, knowing the direction of causal relations is also useful. For example, knowing whether a sad parent caused the marriage to deteriorate or whether the deteriorating marriage caused a parent to become sad can impact how a child responds to family problems.

Much if not most of real-world causal learning involves scenarios that are temporal—the causal events unfold over time—and often the events of interest constitute a change from one state to the next. For example, the right panel of Fig. 1 shows a classic ABA “reversal” design. Jim’s medication is switched between an anti-anxiety medication and a placebo on alternating weeks. Presumably, Jim’s blood pressure would normally be similar from week to week, so a consistent pattern in his blood pressure that correlates with the medication changes would suggest that anxiety does influence blood pressure. Much of real-world causal reasoning (by normal people, not scientists) involves similar comparisons across time. For example, imagine a child trying to figure out how to use a shower. She may turn various knobs and try to identify the difference in the water before versus after the manipulation. Or a teacher might compare his students’ behavior before versus after starting a new motivation technique with the assumption that the students’ behavior is fairly stable over time so a change that coincides with the implementation of the intervention would be a sign of its effectiveness.

In the opposite extreme, there are some instances of causal learning when the states of the variables are only known once. The left panel in Fig. 1 (Atemporal Network) shows the outline of a randomized controlled study in which one group is assigned to an active drug and another is assigned to a placebo, and the outcomes of the two groups are compared. Of course, there is also a middle ground of a pre-post experimental design in which case the experimenters can look at the change within each participant from Time 1 to Time 2 (Fig. 1; Mixed Network). Even though atemporal or “between-subjects” designs are statistically informative, “within-subjects” designs that involve a direct comparison in states over time are more intuitively convincing (Lord & Gilbert, 1983).

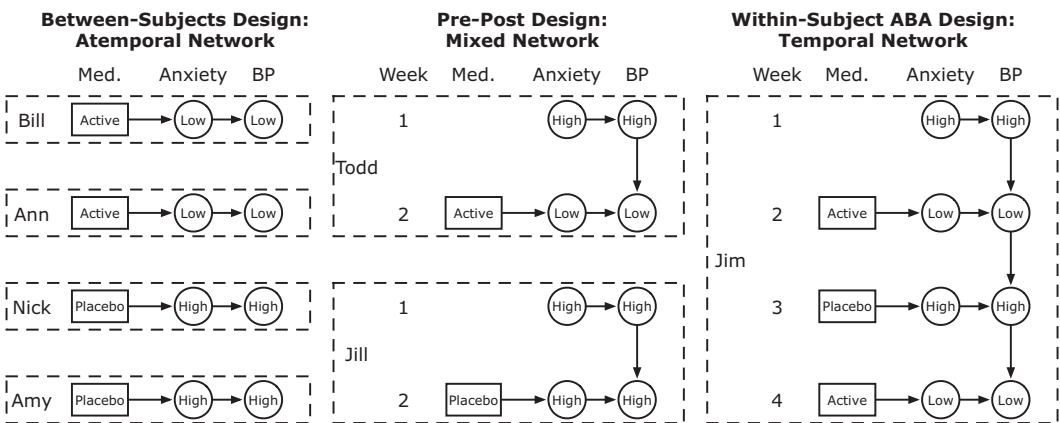


Fig. 1. Three types of experimental designs.

In this study, we examine how children learn the direction of a causal relationship in cases when one causal entity is repeatedly manipulated or observed over time and the state of a variable is influenced by its previous state (temporal dependence or autocorrelation). We introduce Temporal Causal Bayesian Networks to help develop intuitions about temporal causal inference. Specifically, if one changes the state of Variable *X* and sees (vs. does not see) a change in Variable *Y* compared to its prior state, one would likely infer that *X* influences (vs. does not influence) *Y*. Similar intuitions can also facilitate learning the direction of a causal relationship even if one does not have the ability to intervene: If one observes a correlation between *X* and *Y* over time and then sees *Y* change while *X* stays stable, one might infer that *X* probably causes *Y* instead of the reverse. In the following section, we first discuss how causal beliefs can influence causal inference, and then we introduce temporal causal networks to provide a more formal way to capture these intuitions about temporal dependency.

1.1. How causal beliefs influence causal learning

An interesting aspect of causal cognition concerns how adults and even children incorporate many types of beliefs about how a causal mechanism works when interpreting evidence about the functioning of the mechanism. For example, Mendelson and Shultz (1976) presented children with a scenario with two cues that were potential causes of an effect; *Cue A* always occurred 5 s before the effect, and *Cue B* sometimes occurred immediately before the effect, but sometimes the effect occurred in the absence of *Cue B*. If the children had reason to believe the mechanism would have a delay (a marble had to pass through a long chute), they attributed the effect to *A*, but otherwise they attributed the effect to *B*. Lucas, Gopnik, and Griffiths (2010) showed children blocks, some of which could cause a machine to activate when placed on machine. Children in the “conjunctive” condition saw that the machine would activate only if Block *A* and Block *C* were simultaneously put on the machine. In contrast, children in the “disjunctive” condition learned that *A* and *C* could each individually cause the machine to activate. Then, the children saw trials in which *D* failed to activate the machine alone but the combination of *D* and *F* succeeded. Children in the disjunctive condition concluded that *D* was not causally effective (the activation was due to *F*), but children in the conjunctive condition were more likely to conclude that *D* had causal power (in combination with *F*). In summary, beliefs about how a causal mechanism works can influence how one interprets data and draws conclusions.

In the last two decades, the causal Bayesian Network framework (henceforth “causal networks”) has been developed by statisticians and computer scientists to represent and reason about causal relations (e.g., Pearl, 2000; Spirtes, Glymour, & Scheines, 1993), and it has also been proposed as a framework for how both adults (e.g., Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003) and children (e.g., Gopnik, 2004; Gopnik et al., 2004) learn causal relations. One of the benefits of causal networks is that they provide a framework for how different types of beliefs about a causal system can affect inferences about the causal system (e.g., Griffiths & Tenenbaum, 2009; also see Gopnik & Wellman, 2012).

Building off this framework, here we use causal networks to represent beliefs about temporal dependency (i.e., autocorrelation). We are not implying that children necessarily use Causal Bayesian Networks for learning causal relations. Rather, we are using causal structure diagrams as an intuitive and precise way to demonstrate reasoning patterns with temporally dependent causes.

1.2. Causal networks

Fig. 2 shows four causal networks. In all of them, there are two variables, *A* and *B*, each of which can be in the state 0 or 1. Here, we are assuming that the arrows represent positive causal relations. Interventions, external manipulations that set one of the variables to a particular state such as taking an anti-anxiety medication in Fig. 1, are represented with square nodes. Because these interventions completely determine the state (1 or 0) of the manipulated variable, all other arrows representing other causal relationships influencing the manipulated variable are removed. For example, in the structures in which *A* causes *B*, when there is an intervention on *B* the $A \rightarrow B$ arrow is removed, but when there is an intervention on *A* the $A \rightarrow B$ arrow remains intact.

First, consider the atemporal network in Fig. 2. Atemporal causal networks represent scenarios when each variable is only measured once within each instance, analogous to the between-subjects design in Fig. 1. Across all the instances the causal strength of *A* causing *B* is assumed to be the same. However, each instance is independent of the others; the fact that $a = 1$ for Instance 1 has no bearing on the state of *A* for Instance 2.

Now consider the temporal networks in Fig. 2. These networks represent one instance (e.g., one person) over five time periods, analogous to the within-subjects design in Fig. 1. The three networks represent *A* causes *B*, there is no relation between *A* and *B*, and both *A* and *B* influence each other. The critical difference between atemporal and

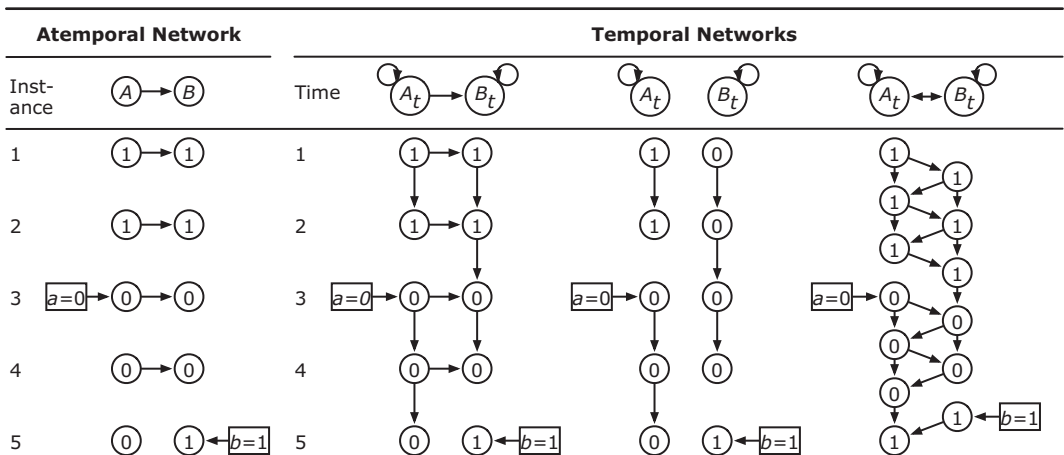


Fig. 2. Example atemporal and temporal causal networks.

temporal networks is that in temporal networks the prior state of a variable can carry forward and influence the current state; all else being equal, variables tend to stay in the same state over time, which is known as positive autocorrelation. The autocorrelation is represented by the downward arrows. The variables in these examples and in the experiments are strongly autocorrelated. But even though the downward arrows are fairly strong, the horizontal relationships are even stronger. For example, at Time 3 in the $A \rightarrow B$ network, the intervention setting $a = 0$ carries over to B even though $b = 1$ at Time 2.

Just as in the atemporal network, the interventions in the temporal networks are deterministic, so any arrows going into the manipulated variable, be they horizontal or vertical, are removed. Also, note how temporal networks can represent bidirectional causal relationships such as $A \leftrightarrow B$. Temporal networks accomplish this by “unfolding” the relationships over time so that at Time 1 A causes B and at Time 2 B causes A .

In summary, the purpose of these figures is to show how causal networks can represent scenarios with temporally dependent or independent variables. A full representation of these causal networks would also include parameters to represent the probability of each node being in a particular state given the states of the nodes that directly cause it (see Rottman & Hastie, 2013, for an introduction to causal networks). But here we focus on intuitions involving the causal graphs alone. In the next section, we demonstrate how these graphs can be used to identify which structure is most likely to have produced a given set of data. The temporal and atemporal graphs sometimes identify different causal structures.

1.3. Learning causal direction from interventions

A number of studies have investigated how children learn causal relations when they have the opportunity to intervene on a causal system (e.g., Gopnik et al., 2004; Kushnir, Gopnik, Lucas, & Schulz, 2010; Schulz, Gopnik, & Glymour, 2007). These studies have demonstrated that preschool children understand principles such that a correlation between two variables implies some causal relationship, that an intervention on A that succeeds in producing an effect B implies a causal relation from A to B . They even infer an unobserved common cause if A and B are correlated, but interventions on either one do not produce the other.

One aspect of these studies is that the events were punctate—they occurred temporarily—but between each trial the variables were in one default “off” state. This meant that the state of the variables in one trial did not impact the state of the variables on subsequent trials. Our goal in this manuscript was to investigate causal learning when variables are temporally extended and dependent—variables remain in states for periods of time and can stay in the same state across trials—which occurs in a vast array of real-world cases. To test this use of information, we use sets of data like the one in Table 1.

Table 1 shows two variables over five time periods. Bold represents a manipulation setting the bolded variable to either 1 or 0. The top part of Fig. 3 (“Temporal Causal Structure Hypotheses”) demonstrates the type of reasoning process that could be used to learn which causal structure is most likely to have generated the data in Table 1

Table 1
Example intervention data

| Time | Order 1 | | Order 2 | |
|---------|----------|----------|----------|----------|
| | <i>X</i> | <i>Y</i> | <i>X</i> | <i>Y</i> |
| Initial | 0 | 0 | 0 | 0 |
| 1 | 1 | 1 | 1 | 1 |
| 2 | 1 | 0 | 0 | 0 |
| 3 | 1 | 1 | 1 | 1 |
| 4 | 0 | 0 | 1 | 0 |

Note. Bold represents a manipulation setting the bolded variable to either 1 or 0.

assuming that the variables are autocorrelated. We will first start by examining the Order 1 data.

The Order 1 data fit well with the Temporal $X \rightarrow Y$ structure. Interventions on X transfer to Y , and when there is an intervention on Y , X stays in its previous state. However, the other structures fit the data worse. At Time 1 X is set to 1, and Y also changes to 1. For the $X \leftarrow Y$ structure and the structure in which X and Y are unrelated this simultaneous change is a coincidence, making these structures less likely. For the $X \leftarrow Y$ and $X \leftrightarrow Y$ structures, when Y is set to 0 at Time 2, X fails to change, making these structures less likely.

Believing that the variables are temporally dependent leads to an interesting implication about the informativeness of interventions. At Time 3, Y is set to 1. But because $x = 1$ at Time 2 and because we are assuming positive causal relationships, the intervention on Y is uninformative; all of the causal structures predict that $x = 1$ at Time 3. This is essentially a “ceiling” effect. Understanding the informativeness of an intervention for a given scenario is critical for causal learning more generally. For example, Cook, Goodman, & Schulz (2011) showed how preschoolers test causal factors independently, which is often more informative than testing them simultaneously, and Wu and Cheng (1999; see also Cheng, 1997) found that college students believe ceiling and floor effects to be uninformative for scenarios like the Atemporal Network and the Mixed Network in Fig. 1 (see also Kushnir, Wellman, & Gelman, 2008).

To demonstrate how beliefs about temporal dependence can influence causal inference, consider Order 2 in Table 1. Order 2 has the exact same set of trials as Order 1 but they are arranged differently in time. The middle part of Fig. 3 shows how the four causal structures fit with Order 2. $X \rightarrow Y$ and $X \leftrightarrow Y$ both are likely. Neither $X \rightarrow Y$ nor $X \leftrightarrow Y$ fits the data perfectly but they fit better than $X \leftarrow Y$ and the no link structure. In summary, believing that a scenario involves autocorrelation, which is captured by the Temporal Causal Structures, can lead a person to infer different causal structures given the same set of data ordered in a different way; $X \rightarrow Y$ was most likely for Order 1 but $X \rightarrow Y$ and $X \leftrightarrow Y$ are both likely for Order 2.

Another way to see the influence the belief that the variables are autocorrelated is to consider the causal structure that would be inferred if one believed that the scenario was

temporally independent (no autocorrelation). The bottom part of Fig. 3 fits the same data

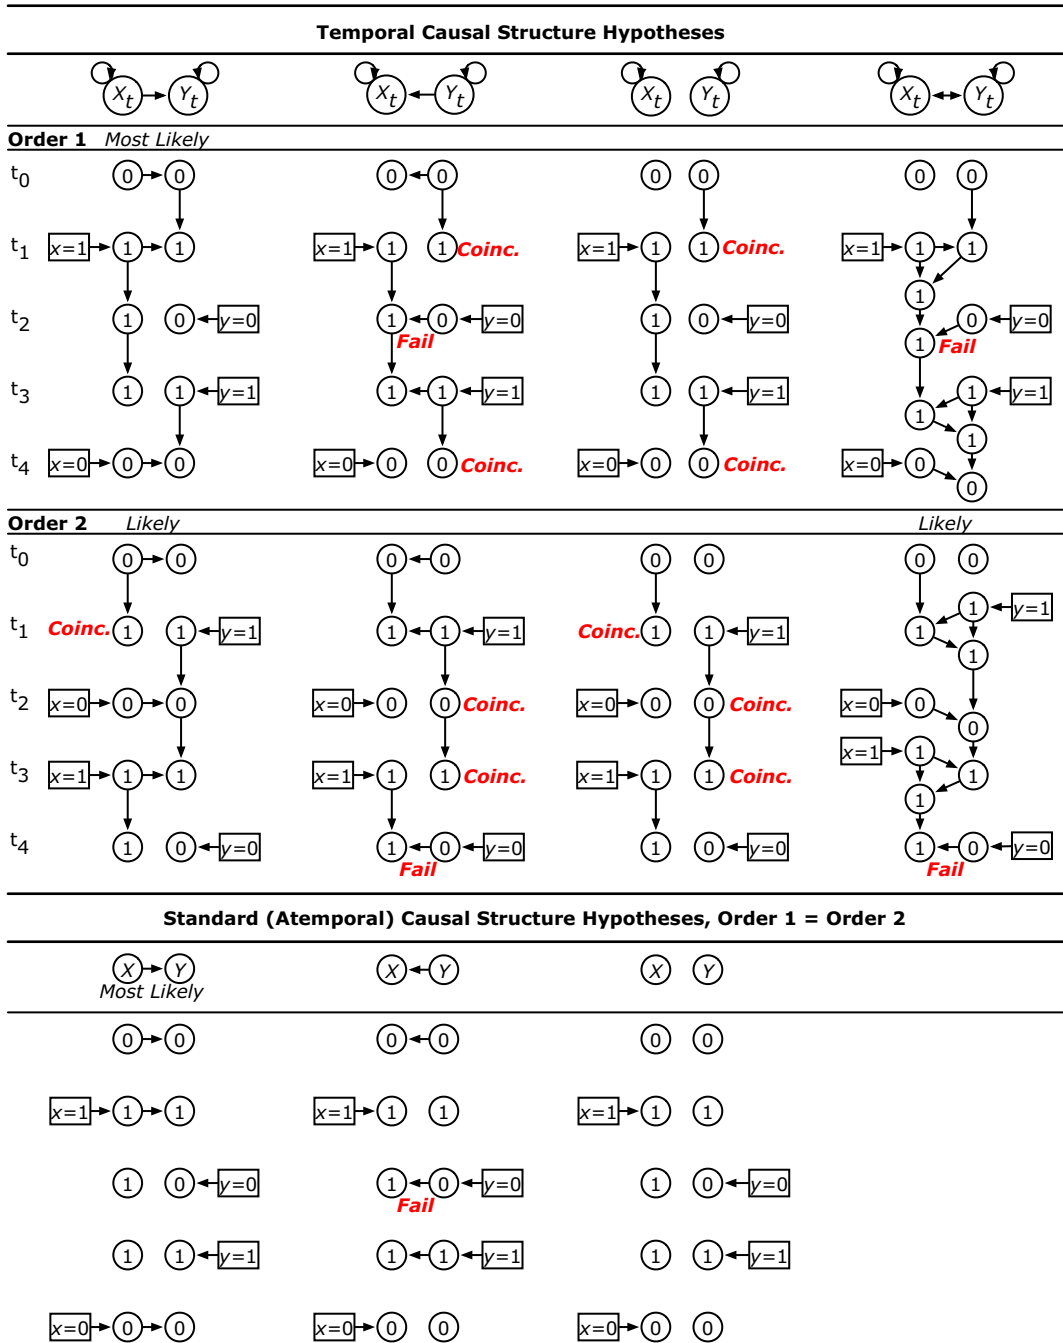


Fig. 3. Fitting temporal and atemporal causal structures to the data in Table 1.

to atemporal causal structures treating each time as a separate instance. Atemporally, there is no difference between the two orders, so we only present Order 1. Both the $X \rightarrow Y$ structure and the no relationship structure are plausible; however, a trend is beginning to appear that interventions on X typically transfer to Y , which would imply $X \rightarrow Y$. The $X \leftarrow Y$ structure is less likely due to the third row when Y was set to 0 but $x = 1$.

In summary, this example demonstrates how one might learn causal directionality from how variables change over time with interventions. Because sometimes the prior states of variables make a particular intervention uninformative, changing the order of the trials can lead to different inferences. In contrast, changing the order of the data does not matter according to atemporal reasoning strategies for which each time period is examined individually.

Finally, we note that Griffiths and Tenenbaum (2009, p. 692–695, see also Griffiths, Baraff, & Tenenbaum, 2004) made a Bayesian model of learning the direction of a causal relationship between two variables (specifically for the “stick machine” from Gopnik et al., 2004, which is also the basis for our experimental setup). Their model incorporates naive beliefs about physical principles about how one element in a machine moves if another one moves, which facilitates inference from a small number of trials. They even made a version of the model for bidirectional causal relations that uses an unfolded temporal causal network. Thus, our model here can be viewed as an extension of theirs for cases with temporal dependency; the previous studies on the stick machine used punctate events.

1.4. Learning causal direction from observations

So far, we have discussed how one might learn the direction of a causal relation from intervening on a temporal causal structure. However, often we do not have the ability to intervene and can merely observe two variables over time. Of course, merely observing a correlation between two variables is insufficient to determine the direction of the causal relation (or if there is a common cause). Temporal priority (e.g., X occurs and then Y subsequently occurs) is a strategy that both children and adults use for inferring causal structure (Burns & McCormack, 2009; Lagnado & Sloman, 2006). Here, we investigate another strategy for learning causal directionality from observations when one believes the variables to be temporally dependent.

The strategy is simple. If Y changes state by itself but X fails to change, this implies that Y does not influence X , or the causal relation from Y to X is fairly weak. X and Y both changing simultaneously implies that there is some relation between X and Y . Further, if both of these two types of changes occur within the set of data, one might infer that X causes Y by process of elimination.

The top of Fig. 4 shows how three temporal causal structures would attempt to explain a set of observational data. In Fig. 4, we introduce a notation that is atypical of causal networks; we include “hypothesized interventions” in the dotted squares. By hypothesized intervention we mean any factor that is hypothesized to be responsible for an observed change. For example, at Time 2 both X and Y change from 1 to 0. If people think about

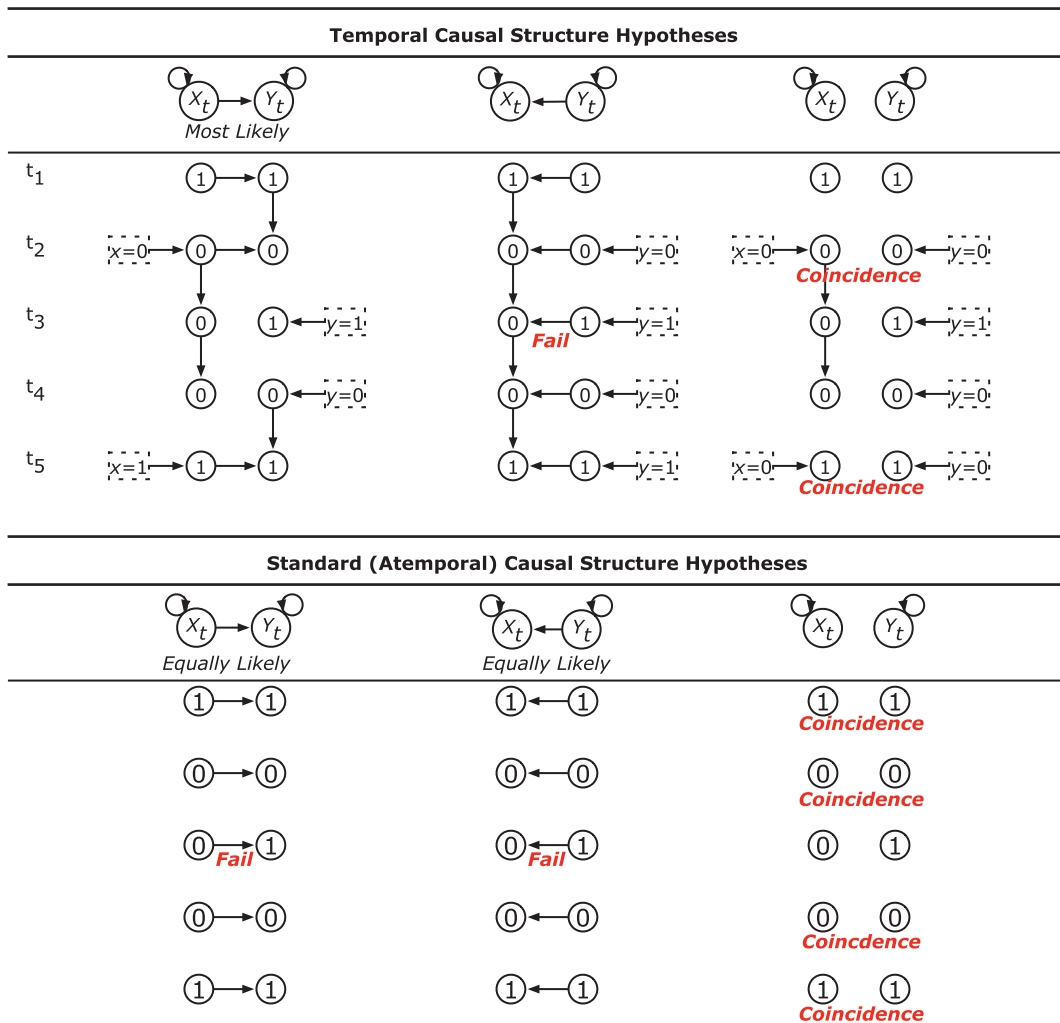


Fig. 4. Three temporal causal structures fit to the same set of data.

possible factors that could have produced this simultaneous change, they could infer that something produced a change in X that carried over to Y , $X \rightarrow Y$, or something produced a change in Y that carried over to X , $X \leftarrow Y$. They could also believe that X and Y are unrelated and that there were external factors that changed both simultaneously, but such coincidental events would probably be viewed as less likely.

At Time 3, Y changes but X does not (and X was not at ceiling). This suggests that Y does not cause X , making structure $X \leftarrow Y$ less likely. This failure is more significant to the extent that one believes that causal relations tend to be fairly strong (e.g., Lu et al., 2008; Yeung & Griffiths, 2011).

There are two other possible temporal causal structures not shown in Fig. 4. A bidirectional causal relationship would also include the same *Fail* as in $X \leftarrow Y$, so is less likely.

Another possibility is that there is a common cause of both X and Y . According to this framework, it would be difficult to distinguish a bidirectional structure with weak links from a common cause; we do not investigate this possibility in the current manuscript.

The bottom part of Fig. 4 shows how atemporal causal networks could be fitted to these data. As researchers, we are all familiar with the notion that a correlation cannot identify a specific causal relation; the bottom part of Fig. 4 shows why visually. The $X \rightarrow Y$ and $X \leftarrow Y$ structures both have one trial when the hypothesized causal relation fails, so they are both equally likely. The no links structure is fairly unlikely because in five of six trials the two variables are in the same state.

In summary, Fig. 4 gives an outline for how temporal reasoning on causal networks can be used to learn the direction of a causal relation even when the two variables are merely observed over time without delay (see Rottman & Keil, 2011; for two fleshed-out quantitative models). More generally, when Figs. 3 and 4 are compared, it can be seen that similar reasoning strategies can be used in interventional and observational learning. The difference is that in the observational case one must hypothesize external factors that produced the changes in the observed variables, which is not necessarily a simple task. For example, at Time 2 in Fig. 4, if one believes that the true structure is $X \rightarrow Y$, a plausible hypothesis is that something produced a change in X , which carried over to Y , though there could have been two external factors that produced simultaneous changes in both X and Y . Alternatively, if one believes that $X \leftarrow Y$ is the most likely structure, then a plausible hypothesis is that something produced a change in Y that carried over to X . Thus, learning causal direction is considerably harder in the observational than interventional case. The key insight is that Y changing without a change in X is evidence that Y does not influence X .

1.5. Motivation for current work and outline

Rottman and Keil (2012) demonstrated that college students use the strategies outlined above for learning causal directionality between two or three variables. They also found that the two strategies are cognitively related; working with an interventional scenario facilitated subsequent use of the observational strategy.

1.5.1. Interventions

Experiment 1 tests whether children ages 3–7 use the temporal-interventional strategy (e.g., Fig. 3) for learning causal structures. The temporal-interventional strategy appears to be simple, so it is plausible that even preschool-aged children would use it; however, it would provide a novel example of how children understand the informativeness of interventions based on the states of the variables before the intervention.

Another motivation for Experiment 1 is that it is unclear whether children used temporal or atemporal strategies in previous research (e.g., Gopnik et al., 2004; Kushnir et al., 2010; see also Schulz et al., 2007). Studies using data sets with punctate variables imply the same causal structure according to atemporal and temporal strategies (see the Online Supplemental Materials for more details). The current experiments use a slightly modified

experimental setup so that variables could remain in the same state over time to disambiguate between atemporal and temporal strategies, which may provide insight into the strategies that children used in previous tasks (see the General Discussion).

1.5.2 Observations

Experiment 2 tests whether children ages 4–7 use the temporal-observational strategy (e.g., Fig. 4). Observational learning is considerably harder than interventional learning, and thus it is unclear whether children would use the observational strategy. Existing research does not clearly identify when the observational strategy might emerge. In the following paragraphs, we discuss the similarities and differences between the inferential capacities required by the temporal-observational strategy and previously studied inferential capacities.

Temporal contiguity, two events occurring close together in time, signifies that they may be causally related. The observational strategy clearly relies upon temporal contiguity. For example, *X* and *Y* changing simultaneously at Time 2 in Fig. 4 implies that there is some causal relation between the two. Children are more likely to infer a causal relationship when there is less temporal delay between two events (e.g., Mendelson & Shultz, 1976; Siegler & Liebert, 1974) and even 7-month-old children are sensitive to subtle aspects of temporal and physical contiguity (Newman, Choi, Wynn, & Scholl, 2008).

A notion related to but distinct from temporal contiguity is temporal priority; causes occur before effects. Children do use temporal priority for inferring causal relations (e.g., Bullock & Gelman, 1979; Burns & McCormack, 2009). However, temporal priority is not necessary for the current observational strategy; in Fig. 4 the two variables change simultaneously.

Finally, the observational strategy requires a particular use of negative evidence; *Y* changing by itself implies that *Y* does not influence *X*. This particular use of negative evidence has not been studied in children. However, it does have parallels to an atemporal-interventional strategy that if an intervention on *Y* fails to produce *X*, *Y* is inferred not to cause *X*, which is used by children at least as young as 4 (Gopnik et al., 2004, p. 27).

In summary, the temporal-observational strategy requires some familiar and some new inferential capacities. Our goal for Experiment 2 is to uncover whether children use this relatively sophisticated strategy for inferring causal direction at all or if it emerges later in development.

2. Experiment 1: Interventions

2.1. Method

2.1.1. Participants

Children of ages 3–7 were either recruited from the local population around New Haven, CT, and participated in the experiment in a laboratory at Yale University or they were recruited from and participated in the experiment at a children's museum in a

nearby town. Sixty-four children participated. Children were rewarded for their participation with a small toy and certificate of appreciation. Families that participated in the study in the laboratory also received passes to a nearby museum.

2.1.2. Stimuli

The children worked with a computer program running on a touch-screen monitor that simulated the stick-machine used by previous researchers (e.g., Gopnik et al., 2004; Kushnir et al., 2010; see Fig. 5). A video of the computer program can be found as Movie S1 in the Supplementary Material file. The stick-machine comprised two sticks that protrude from a box. The sticks can move up or down, together or independently. In previous research with the stick-machine, the sticks would temporarily move up but would always come back down between each trial. We modified the scenario slightly so that the sticks could remain up for periods of time, reflecting temporally extended states rather than punctate events.

Four practice scenarios helped the children become comfortable working with the machine (see Table 2). In Practice Scenario A, whenever Y was pulled up or pushed down, X also rose up or lowered down, implying $X \leftarrow Y$. In Practice Scenario B, whenever X was pulled up or pushed down, Y also rose up or lowered down, implying $X \rightarrow Y$. In Practice Scenario C, whenever either X or Y was pulled up or pushed down, the other stick also rose up or lowered down, implying $X \leftrightarrow Y$. In Practice Scenario D, whenever either X or Y was pulled up or pushed down, the other stick did not move, implying no relationship between X and Y . These four practice scenarios were designed so that both the atemporal and temporal strategies imply the same causal direction.¹ These four practice scenarios were also designed to suggest to the children that either stick could be a cause or effect, or both or neither.

There were two test conditions (see Table 2). Critically, the two test conditions had the exact same set of trials, just in different orders. Thus, according to the atemporal causal network framework, both of the test conditions imply $X \rightarrow Y$. In both, whenever X

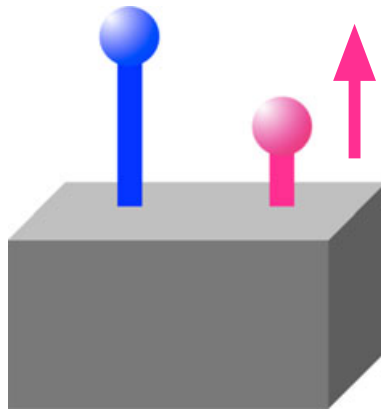


Fig. 5. Sample screenshot from Experiment 1.

Table 2
Stimuli used in Experiment 1

| Time | Practice Scenarios | | | | | | | | Test Scenarios | | | |
|-----------------------------------|--------------------|----------|-------------------|----------|--|----------|----------|----------|-------------------|----------|-----------------------|----------|
| | A | | B | | C | | D | | Order 1 | | Order 2 | |
| | X | Y | X | Y | X | Y | X | Y | X | Y | X | Y |
| Initial | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 0 | 1 |
| 2 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 |
| 3 | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 1 | 1 | 0 | 0 |
| 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
| 5 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 1 | 0 |
| 6 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 7 | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 0 | 1 | 1 | 1 |
| 8 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 9 | | | | | | | | | 1 | 1 | 0 | 1 |
| 10 | | | | | | | | | 1 | 0 | 1 | 1 |
| 11 | | | | | | | | | 1 | 1 | 0 | 0 |
| 12 | | | | | | | | | 0 | 0 | 1 | 1 |
| 13 | | | | | | | | | 1 | 1 | 1 | 0 |
| 14 | | | | | | | | | 0 | 0 | 0 | 0 |
| 15 | | | | | | | | | 0 | 1 | 1 | 1 |
| 16 | | | | | | | | | 0 | 0 | 0 | 0 |
| Predictions of the Two Strategies | | | | | | | | | | | | |
| Temporal Strategy | $X \leftarrow Y$ | | $X \rightarrow Y$ | | $X \leftrightarrow Y$ | | $X; Y$ | | $X \rightarrow Y$ | | $X \leftrightarrow Y$ | |
| Atemporal Strategy | $X \leftarrow Y$ | | $X \rightarrow Y$ | | $X \rightarrow Y$ or $X \leftarrow Y$ | | $X; Y$ | | $X \rightarrow Y$ | | $X \rightarrow Y$ | |

Notes. 0 = stick is down; 1 = stick is up. Boldface represents an intervention pulling the stick up or pushing it down, whereas regular type represents an observation.

was pulled up, Y would be in the up position, and whenever X was pushed down, Y would be in the down position, and no matter whether Y was pulled up or pushed down, X had a 50% chance of being up and a 50% chance of being down.

However, the two orders implied different structures according to the temporal strategy. In the Order 1 condition, whenever Stick X was pulled up or pushed down, Y changed simultaneously. However, whenever Stick Y was pulled up or pushed down, Stick X did not move. The temporal strategy predicts $X \rightarrow Y$. Order 2 is similar to Order 1 for Stick X: whenever X is pulled up or pushed down Y changed simultaneously. However, Order 2 is different from Order 1 for Stick Y: sometimes when Y is pulled up or pushed down X changes simultaneously, and sometimes it does not. Thus, the temporal strategy predicts $X \leftrightarrow Y$ (technically $X \rightarrow Y$ is stronger than $X \leftarrow Y$).

The children first worked with the four practice scenarios in the order A, B, C, D. Then, they worked with the two test conditions in a counterbalanced order. For the two test conditions, the X and Y sticks were counterbalanced to the left and right positions. For each of the six scenarios, the machine and each of the sticks had unique colors emphasizing that the machines were different.

2.1.3. Procedure

When the experiment was conducted at Yale University, the child and experimenter worked together in one room, and parents observed from an adjacent room using a video feed. When the experiment was conducted at the children's museum, parents were allowed to be in the room but were instructed not to react to the experiment and sat outside of the child's field of view.

At the beginning of the experiment, the experimenter started the computer on Practice Scenario A (the pink stick was the cause) and said, "Today, we are going to play a game.... We have a box with two sticks, the blue stick and the pink stick! Some sticks are special, they do something to other sticks. Let's figure out which sticks do something to other sticks!"

On each trial, the computer prompted the child to pull one stick up or push one stick down with an up or down arrow next to the respective stick (see Fig. 5). After the pink up arrow appeared, the experimenter said, "Let's pull the pink stick up! To do that we just touch the stick. Can you touch the pink stick to pull it up? Great work! Oh look; now, the blue stick is up."² Then, the next arrow would appear (e.g., a blue down arrow). This same running dialog was used throughout the first practice scenario. At the end of the first practice scenario, the experimenter said, "The pink stick was special; it did something to the blue stick. The blue stick was not special; it did not do anything to the pink stick."

Then, the children moved on to the second practice scenario. The experimenter said, "Now we are going to play this game another time with another box. This box has a brown stick and a yellow stick. I want you to pay attention and figure out which sticks are special—which sticks do something to the other sticks." For all the remaining practice scenarios and the test scenarios, the experimenter told the child which stick to pull up or push down but did not give the children any feedback about the "correct" answer. At the end of each scenario, the experimenter asked the child whether each stick was "special: does it do something to the [the color of the other stick, e.g., blue] stick?"³

2.2. Results

The atemporal strategy predicts $X \rightarrow Y$ in both conditions. In contrast, the temporal strategy predicts $X \rightarrow Y$ for the Order 1 condition and $X \leftrightarrow Y$ for the Order 2 condition. Fig. 6 shows the percent of participants who endorsed each causal relation. In Fig. 6, we grouped the participants by age for visualization only; in the statistical tests, we used age in days and tested for trends, not differences in specific years.

We ran separate logistic regressions for each causal relation, with age and condition as predictors, and random effects by participant to account for the repeated-measures. For the $X \rightarrow Y$ link, there was a main effect of age; older children were more likely to endorse this link than younger children, $p = .01$ (all ps reported in this manuscript are two-tailed). As both strategies predict endorsement of the $X \rightarrow Y$ link, this developmental pattern probably reflects both a developing ability to infer causal relations as well as developing general cognitive abilities. There was no effect of condition nor an interaction, $ps > .38$.

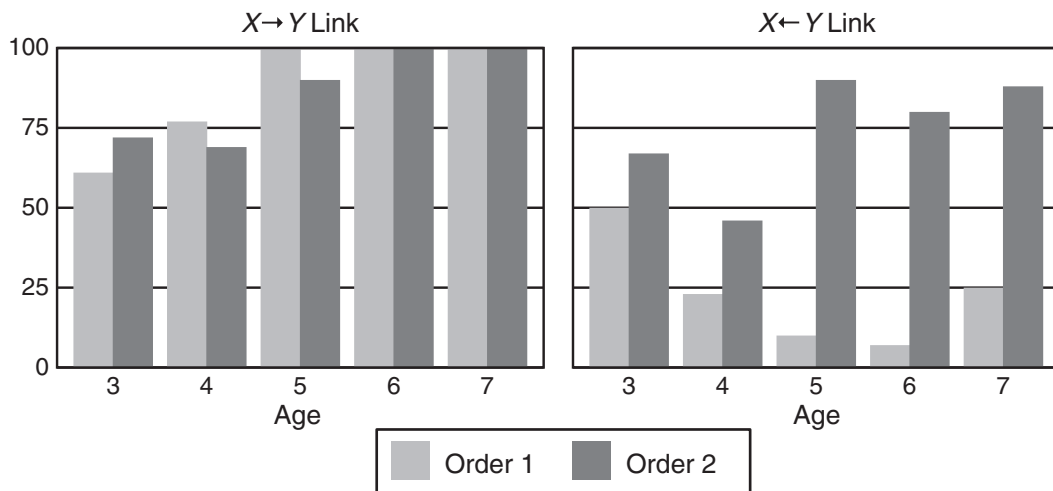


Fig. 6. Percent of participants who inferred each causal relation by age and condition.

The critical inference is the $X \leftarrow Y$ link. A logistic regression found a significant effect of condition, $p < .01$. As can be seen in Fig. 6, participants were more likely to endorse the $X \leftarrow Y$ link in the Order 2 condition than the Order 1 condition. There was no main effect of age, $p = .10$, but there was a significant interaction between condition and age, $p < .01$; the difference between the two conditions gets larger with age, which likely reflects more use of the temporal strategy. Follow-up tests of the two conditions separately revealed that endorsement of the $X \leftarrow Y$ link decreased with age in the Order 1 condition, $p = .02$, and increased (marginally) in the Order 2 condition, $p = .06$.

While the difference for endorsing $X \leftarrow Y$ across the two conditions is smaller in the younger children, this result does not imply use of the atemporal strategy over the temporal strategy. The atemporal strategy predicts no endorsement of the $X \leftarrow Y$ link in either condition. In contrast, the younger children's inferences are basically at chance (50%) responding.

In summary, children's inferences are more consistent with the temporal strategy than the atemporal strategy, and their use of the temporal strategy increases with age.

3. Experiment 2: Learning causal direction from observations over time

In Experiment 2, we tested whether children would use the temporal-observational strategy outlined in the introduction for learning the direction of a causal relation when one believes that the variables are temporally dependent.

3.1. Methods

3.1.1. Participants

Forty-two children aged 4–7 were recruited from the same populations as in Experiment 1.

3.1.2. *Stimuli and design*

Children first worked with two interventional training conditions and then four observational test conditions. In the training conditions, the children observed a cartoon monkey inside the stick-machine manipulating the sticks. One purpose of the interventional training conditions was to help the children understand the stick-machine; that one stick can affect the other stick.

The two training conditions used the same set of data as Practice Scenarios A and B from Experiment 1 (see Table 2). The main difference from Experiment 1 was that the monkey, not the children, manipulated the sticks. The atemporal and temporal strategies imply the same causal direction for these training scenarios.

In the test conditions, the children were told that the monkey was inside the stick-machine, but the children could no longer see the monkey performing the interventions; they could only see the sticks moving. Again the goal was to figure out which stick influenced the other one.⁴ Having the monkey inside the machine provided the children with a concrete explanation for the sticks' movements (e.g., the pink stick moved up; the monkey must have pushed it), lest the movements be totally unexplainable.

The test conditions used the data sets in Table 3. According to the atemporal strategy, there is a correlation between the two sticks of .5; however, this correlation does not imply a causal direction. However, looking at how the sticks change over time, the Y stick sometimes changed position on its own (which implies that Y does not influence X),

Table 3
Summary of stimuli in test conditions in Experiment 2

| Time | Test A | | Test B | |
|---------|--------|---|--------|---|
| | X | Y | X | Y |
| Initial | 0 | 0 | 0 | 0 |
| 1 | 1 | 1 | 0 | 1 |
| 2 | 1 | 0 | 0 | 0 |
| 3 | 1 | 1 | 1 | 1 |
| 4 | 0 | 0 | 1 | 0 |
| 5 | 1 | 1 | 1 | 1 |
| 6 | 0 | 0 | 0 | 0 |
| 7 | 0 | 1 | 1 | 1 |
| 8 | 0 | 0 | 0 | 0 |
| 9 | 0 | 1 | 1 | 1 |
| 10 | 0 | 0 | 1 | 0 |
| 11 | 1 | 1 | 1 | 1 |
| 12 | 0 | 0 | 0 | 0 |
| 13 | 1 | 1 | 0 | 1 |
| 14 | 1 | 0 | 0 | 0 |
| 15 | 1 | 1 | 1 | 1 |
| 16 | 0 | 0 | 0 | 0 |

Note. 0 = stick is down; 1 = stick is up.

and sometimes both X and Y changed position together. According to the temporal reasoning strategy explained above this pattern of changes implies $X \rightarrow Y$.

Overall, there were four test datasets (Table 3). Two more were created from the Test A and Test B conditions by switching Trials 1–8 and Trials 9–16 within the conditions, respectively. The X and Y sticks in Table 3 were counterbalanced to the left and right positions.

3.1.3. Procedure

A video of this experiment can be found as Movie S2 in the Supplementary Material file. The children were first introduced to the training scenarios. They were told that “We are going to play a game where you watch a monkey play with a few new toys.... The monkey is in a box with two sticks, the black stick and the gray stick! Some sticks are special; they do something to other sticks. Let’s figure out which sticks do something to other sticks! When we press the bell, it tells the monkey to move a stick. The monkey gets to decide which stick he wants to move. Try pressing on the bell! Great! Look, the monkey moved the gray stick and the black stick is up!” (see Fig. 7 for a screenshot). Between each trial, the children rang the bell to cue the monkey to manipulate a stick. The monkey only moved one stick at a time. In addition, the monkey only “intervened” to push a stick up or “pull” one down, not to hold a stick in place, same as in Experiment 1.

At the end of the first training scenario, the children were told that “The gray stick was special; it did something to the black stick! The black stick was not special; it did not do anything to the gray stick.” The children were not given feedback about which stick was “special” for the second training example.

After the two training examples the children were introduced to the testing conditions. They were told “OK, now you are going to see a few more toys. You won’t be able to

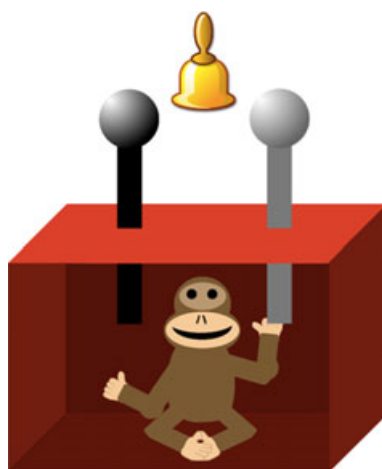


Fig. 7. Screenshot of a training condition in Experiment 2 in which the monkey is in the process of pushing the stick on the right up and the stick on the left is simultaneously rising.

see the monkey, but you will use the bell to tell him when you want him to move a stick. Just like before, I want you to figure out which sticks are special, so which sticks do something to other sticks.” The children saw a machine like the one in Fig. 7, only the front of the machine was opaque so that they could not see the monkey. During each of the testing conditions, to keep the children’s attention, the children were periodically asked, “How many sticks moved?” At the end of each testing scenario, the experimenter asked the child whether each stick was “special: does it do something to the [the color of the other stick, e.g., gray] stick?”

In piloting, it was found that some children lost interest in the task after two test scenarios. To ensure that the children were engaged in the task for every scenario they completed, after the second test scenario they were informed that there were two additional toys (test scenarios) they could play with if they were interested, or they could choose to stop.

3.2. Results and discussion

The atemporal strategy does not imply one direction over the other. In contrast, the temporal strategy implies that $X \rightarrow Y$ (see Table 3 and Fig. 8).

Participants usually inferred that $X \rightarrow Y$ or $X \leftarrow Y$, not a bidirectional link or no link. This is likely due to the two training conditions; one was intended to imply $X \rightarrow Y$ and the other $X \leftarrow Y$. Because these dependent measures were correlated ($r = -.65, p < .01$), we first analyzed the two links separately.

To test whether participants were more likely than chance (.5) to endorse the $X \rightarrow Y$ link, and less likely than chance to endorse the $X \leftarrow Y$ link, we used logistic regressions with no predictors (just an intercept) and random effects to account for the repeated measures within each participant. Both of these predictions were supported by the model,

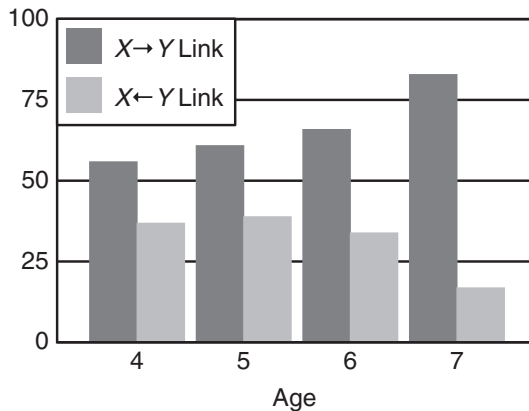


Fig. 8. Percent of participants who inferred each causal relation by age. Note. Ages were binned by year for visualization only.

with the following means (the proportion of endorsing a given link) and 95% confidence intervals: $M_{X \rightarrow Y} = .70$, $CI_{X \rightarrow Y} = (.57-.75)$ and $M_{X \leftarrow Y} = .31$, $CI_{X \leftarrow Y} = (.22-.41)$, $ps < .01$.

Second, we added age into both of these logistic regressions. Even though endorsement of $X \rightarrow Y$ appears to increase with age in Fig. 8 ($p = .12$), and endorsement of $X \leftarrow Y$ appears to decrease ($p = .34$), these relationships were not statistically significant.

We also analyzed the data putting both inferences into the same logistic regression. There was a significant difference in the probabilities of endorsing the two links, $p < .01$. There was a marginal interaction between age and direction of link, $p = .09$, which reflects the increasing difference between endorsing $X \rightarrow Y$ over $X \leftarrow Y$ over age in Fig. 8.

In summary, this experiment found that children are able to infer causal direction based on how two variables change over time. Though none of the tests found statistically significant effects of age, it is likely that a wider sample of ages would find a more robust developmental trend in the use of this strategy for inferring causal direction.

4. General discussion

We examined the strategies that children use to infer the direction of causal relations in scenarios when the variables are autocorrelated. In the first experiment, we compared two conditions. In one condition, whenever a variable X was manipulated, another variable Y changed compared to its prior state, but when Y was manipulated, X stayed stable in its prior state; children primarily inferred $X \rightarrow Y$. In a second condition, sometimes when either X or Y was manipulated, the other variable changed compared to its prior state; children primarily inferred $X \leftrightarrow Y$. Critically, atemporal strategies which only look at the states of the variables after an intervention (not change compared to before an intervention) would yield the same inferences in both conditions.

In the second experiment, we examined how children inferred causal direction when they could only observe two variables over time, not intervene. The atemporal strategy only identifies a correlation between the two variables. However, tracking the variables over time revealed the following pattern: When X changed state, Y also changed state, but sometimes Y changed state without X changing state. Since X sometimes remained stable when Y changed, children were more likely to infer $X \rightarrow Y$.

4.1. *The ambiguity of temporal reasoning in previous studies*

The current experiments sought to identify the strategy that children use for learning causal direction when working with one causal mechanism over time. In this context, children apparently compare the states of variables to their prior states, which makes sense if they believe variables to be autocorrelated.

Furthermore, it is possible that these same temporal comparison strategies (comparing the current state to the state immediately before it) may have been used in previous experiments (see the Online Supplemental Material for details). For example, the Blicket Detector tasks involve placing objects on a detector and seeing which objects activate the

detector. Typically, each trial is treated as independent from the previous trial and the proposed reasoning strategies involve calculating summary statistics such as conditional probabilities on the entire set of trials (e.g., *A* activates the detector by itself, but *B* does not activate the detector by itself, only when *A* is also on the detector). However, it is entirely possible that the participants in those tasks reasoned about how putting an object on the detector or taking it off changed the activation of the detector compared to before the intervention, and same for the previous studies using the stick-machine. It is not possible to clearly disambiguate use of temporal and atemporal strategies in those studies.

In summary, in many situations atemporal and temporal strategies imply the same causal structure. To our knowledge, there has not been a test of how children learn causal structure in a scenario in which the trials are clearly and unambiguously framed as independent, which is the type of scenario for which the atemporal causal network framework is intended to apply. An example of such a scenario is the between-subjects design in Fig. 1. We acknowledge that designing such a test is challenging because even if a scenario is framed with independent events (like a controlled randomized trial), often the learning data are displayed sequentially, potentially introducing the possibility of temporal factors; perhaps, a summary format is ideal. At the current stage, it is simply unknown whether children use the atemporal strategy and how well. The clearest case of atemporal causal structure learning in adults in which the trials were probably viewed by participants to be independent revealed some, but far from perfect or universal use of the strategy (Steyvers et al., 2003).

4.2. *A list of temporal cues to causality*

A wide range of cues can be used to infer causality, including several temporal cues. We suspect that often many of these cues are used together in a converging fashion and that some may often be early emerging in development. There remains a clear need to better understand how each of these cues interact over the course of development and how they fit with other atemporal cues.

Unfortunately, however, the research community does not yet have a comprehensive taxonomy of all the temporal cues to causality (see also Schultz, Fisher, Pratt, & Rulf, 1986). Further complicating the efforts to build a taxonomy is that these cues can sometimes be hard to distinguish in various paradigms and they are not mutually exclusive. We already discussed four cues that can be used to infer causal direction: the temporal-interventional strategy, the temporal-observational strategy, temporal contiguity, and temporal primacy. Here, we list some other temporal cues to causality, though they are not specifically focused on inferring causal direction.

4.2.1. *Primacy/recency*

The order of a series of events often plays an important role in causal inference. For example, if a causal relation is first judged to be very strong, people often discount subsequent evidence that it is weak. However, sometimes later evidence influences the

final judgment more than initial evidence (e.g., Dennis & Ahn, 2001; see also Fernbach & Sloman, 2009; Luhmann & Ahn, 2007).

4.2.2. *Inferring time periods or unobserved factors*

If one notices that a causal relation behaves consistently for a period of time and later starts behaving in a different way, one might infer that some unobserved factor changed (e.g., Buchanan & Sobel, 2011; Gershman, Blei, & Niv, 2010; Redish, Jensen, Johnson, & Kurth-Nelson, 2007; Rottman & Ahn, 2011). Believing that something about the world has changed but that the new state of the world is relatively stable is one rational reason for recency effects; experiences farther away in time are less informative of the current functioning of the world.

4.2.3. *Other cues*

There are many other cues such as whether one variable influences the rate (e.g., Griffiths & Tenenbaum, 2005) or duration of another variable. In the animal learning literature, there many other experimental paradigms such as preconditioning and occasion setting, many of which have natural “causal” interpretations (e.g., many studies with children have investigated scenarios similar to blocking and retrospective revaluation; see Gopnik, Sobel, Schulz, & Glymour, 2001, for a discussion).

4.3. *Causal sufficiency and broader theoretical considerations*

Another way to learn the direction of a causal relation from merely observing two variables occurs when people have a prior belief with causes that are sufficient to produce effects (Mayrhofer & Waldmann, 2011). If one believes that, in a particular domain, causes are sufficient to produce their effects (i.e., $x = 1$ is sufficient to produce $y = 1$), and one observes that whenever $x = 1$, $y = 1$, but sometimes $y = 1$ even if $x = 0$ (i.e., $y = 1$ is not sufficient to produce $x = 1$), then one would infer that X causes Y and not vice versa. The temporal cue for Experiment 2 is somewhat similar; when the state of a cause is changed, it is believed to be sufficient to produce a change in its effect. Thus, if Y changes but X stays stable, one would infer that Y does not cause X . However, these strategies are distinct in that one applies to an atemporal scenario (the states of the variables) and one applies to a temporal scenario (how the states of variables change over time). The atemporal version of the causal sufficiency heuristic cannot explain the current results because $x = 1$ was not sufficient or necessary for $y = 1$.

The similarity between Mayrhofer and Waldmann’s (2011) sufficiency heuristic and the temporal strategy for observations tested here raises a broader question about the similarity between temporal and atemporal strategies. One way to think about the relationships between temporal and atemporal strategies is with an analogy to t -tests. Within-subjects t -tests can be viewed as “temporal”—they look at the change within one entity (or otherwise correlated data). Between-subjects t -tests and one-sample t -tests are “atemporal.” Of course, if one takes the difference between the two time periods, one can

use a one-sample *t*-test instead of a within-subjects *t*-test; in some instances, “temporal” strategies are mathematically equivalent, given the right comparison in data points, to atemporal strategies.

However, the important psychological aspect of these inferences is how they are performed, what pieces of information are being compared, and how the learner conceives of the scenario. The current experiments demonstrate that children conceive of these scenarios as temporal, and they are using temporal comparisons to perform the inference. It is possible to think of these scenarios as either temporal or atemporal, and an important further question is to investigate the cues that children use to interpret the scenario, which is analogous to how a statistician decides which statistical test is correct for a given situation. We previously found that adults pick up on and use temporal cues even in situations when the cover story frames trials as independent (Rottman & Keil, 2012); so, it is likely that the temporal strategy is hard to “turn off.” Because the temporal strategy may be difficult to suppress even when such suppression is warranted, it may be the default option in most naturalistic situations, a default that seems to be present in young children. If this default interpretation is correct, it may be that younger children have even more difficulty than older children in adults in “turning off” the default strategy when the situation clearly indicates that trials are completely independent.

4.4. Conclusions

Beliefs about the direction of a causal relation are critical for deciding which variable to manipulate to produce a desired outcome and for explaining why a particular event occurred. More broadly, people often reason about systems of causal relations between multiple variables. Learning the direction of a relation between two variables likely serves as the fundamental building block for learning whole systems of causal relations (e.g., Waldmann, Cheng, Hagmayer, & Blaisdell, 2008; Fernbach & Sloman, 2009, Experiment 2). The current research demonstrates that children use changes in variables over time to infer the direction of a causal relation. In the future, it will be important to identify how specific contexts elicit the use of particular strategies for learning causal structures and whether children use the various strategies in appropriate and adaptive ways.

Acknowledgments

The authors thank the staff and guests of Stepping Stones Children’s Museum in Norwalk, CT, for their contributions to the studies reported here. This research was supported by NIH Grant F32 HL108711 (Rottman) and NIH Grant R37 HD023922 (Keil).

Notes

1. Technically, because the atemporal causal network framework does not include $X \leftrightarrow Y$ as a possible network, $X \rightarrow Y$ and $X \leftarrow Y$ would be equally likely for Practice Scenario C.
2. These statements were intended to call the attention of the children to the other stick. We phrased them to call the attention to the position of the other stick (atemporal strategy), not whether the other stick changed position (temporal strategy).
3. Gopnik et al. (2004) defined “special” for the children as making the other stick “move.” We chose to use the more abstract “do something” because we were worried that Gopnik et al.’s language could encourage a temporal strategy—when Y is pulled up, it makes X change position from the previous trial, as opposed to when Y is pulled up X also is up (regardless of its prior position). We acknowledge that even “do” is not a perfect term, but most causal verbs imply a temporal change as opposed to merely a statistical dependency.
4. One possible complication of this experiment is that because children believe that there is a monkey controlling the sticks, the experiment actually involves three factors, not two. We could represent the structure as $Monkey \rightarrow X \rightarrow Y \leftarrow Monkey$, also technically with a higher order link representing the fact that the monkey only intervenes on one stick at a time. Critically, however, the monkey’s actions are never observed. In real-world causal inference, there are always unobserved or unknown factors that produce changes in the observed variables, so we do not believe that the monkey fundamentally alters the scenario.

References

- Buchanan, D. W., & Sobel, D. M. (2011). Children posit hidden causes to explain causal variability. In L. Carlson, C. Hölscher & T. Shipley (Eds.), *Proceedings of the 33rd annual meeting of the cognitive science society* (pp. 3098–3101). Austin, TX: Cognitive Science Society.
- Bullock, M., & Gelman, R. (1979). Preschool children’s assumptions about cause and effect: Temporal Ordering. *Child Development*, 50, 89–96.
- Burns, P., & McCormack, T. (2009). Temporal information and children’s and adults’ causal inferences. *Thinking & Reasoning*, 15(2), 167–196.
- Cheng, P. W. (1997). From covariation to causation: a causal power theory. *Psychological Review*, 104(2), 367–405.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: spontaneous experiments in preschoolers’ exploratory play. *Cognition*, 120(3), 341–349.
- Dennis, M. J., & Ahn, W. (2001). Primacy in causal strength judgments. *Memory & Cognition*, 29, 152–164.
- Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 678–693.
- Forsyth, C. B., Shannon, K. M., Kordower, J. H., Voigt, R. M., Shaikh, M., Jaglin, J. a, Estes, J. D., Dodiya, H. B. & Keshavarzian, A. (2011). Increased intestinal permeability correlates with sigmoid mucosa alpha-Synuclein staining and endotoxin exposure markers in early Parkinson’s disease. *PLoS ONE*, 6(12), e28032. doi:10.1371/journal.pone.0028032.

- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*(1), 197–209.
- Gopnik, A. (2004). Mechanisms of theory formation in young children. *TRENDS in Cognitive Sciences*, *8*(8), 71–377.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*(1), 3–32.
- 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*(5), 620–629.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, *138*(6), 1085–1108.
- Griffiths, T. L., Baraff, E. R., & Tenenbaum, J. B. (2004). Using physical theories to infer hidden causal structure. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society* (pp. 500–505). Austin, TX: Cognitive Science Society.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–384.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*(4), 661–716.
- Kushnir, T., Gopnik, A., Lucas, C., & Schulz, L. (2010). Inferring hidden causal structure. *Cognitive science*, *34*(1), 148–160.
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition*, *107*, 1084–1092.
- Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 451–460.
- Lord, C. G., & Gilbert, D. T. (1983). The “same person” heuristic: an attributional procedure based on an assumption about person similarity. *Journal of Personality and Social Psychology*, *45*(4), 751–762.
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological review*, *115*(4), 955–984.
- Lucas, C. G., Gopnik, A., & Griffiths, T. L. (2010). Developmental differences in learning the forms of causal relationships. In R. Camtrabone & S. Ohlsson (Eds.), *Proceedings of the 32nd annual conference of the cognitive science society* (pp. 2852–2857). Austin, TX: Cognitive Science Society.
- Luhmann, C. C., & Ahn, W. (2007). BUCKLE: A model of unobserved cause learning. *Psychological Review*, *114*(3), 657–677.
- Mayrhofer, R., & Waldmann, M. R. (2011). Heuristics in covariation-based induction of causal models: Sufficiency and necessity priors. In C. H. Carlson & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 3110–3115). Austin, TX: Cognitive Science Society.
- Mendelson, R., & Shultz, T. R. (1976). Covariation and temporal contiguity as principles of causal inference in young children. *Journal of Experimental Child Psychology*, *22*, 408–412.
- Newman, G. E., Choi, H., Wynn, K., & Scholl, B. J. (2008). The origins of causal perception: Evidence from postictive processing in infancy. *Cognitive Psychology*, *57*, 262–291.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
- Redish, A. D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: Implications for addiction, relapse, and problem gambling. *Psychological Review*, *114*(3), 784–805.
- Rottman, B. M., & Ahn, W. (2011). Effect of grouping of evidence types on learning about interactions between observed and unobserved causes. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *37*(6), 1432–1448.
- Rottman, B. M., & Hastie, R. (2013). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, Advance online publication. doi:10.1037/a0031903.

- Rottman, B. M., & Keil, F. C. (2011). Learning causal direction from repeated observations over time. In L. Carlson, C. Hölscher & T. Shipley (Eds.), *Proceedings of the 33th annual conference of the cognitive science society*. (pp. 1847–1852). Austin, TX: Cognitive Science Society.
- Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*, 93–125.
- Schultz, T. R., Fisher, G. W., Pratt, C. C., & Rulf, S. (1986). Selection of causal rules. *Child Development*, *57*, 143–152.
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental science*, *10*(3), 322–332.
- Siegler, R. S., & Liebert, R. M. (1974). Effects of contiguity, regularity, and age on children's inferences. *Developmental Psychology*, *10*, 574–579.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. New York: Springer-Verlag.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind. Prospects for Bayesian cognitive science* (pp. 453–484). Oxford, UK: Oxford University Press.
- White, J. F. (2003). Intestinal pathophysiology in Autism. *Experimental Biology and Medicine*, *228*, 639–649.
- Wu, M., & Cheng, P. W. (1999). Why causation need not follow from statistical association: boundary conditions for the evaluation of generative and preventive causal powers. *Psychological Science*, *10*(2), 92–97.
- Yeung, S., & Griffiths, T. L. (2011). Estimating human priors on causal strength. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 1709–1714). Austin, TX: Cognitive Science Society.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Movie S1. Movie of Experiment 1

Movie S2. Movie of Experiment 2

Data S1. Comparison of Temporal vs. Atemporal Strategies in Previous Studies