

# Connecting Causal Events: Learning Causal Structures Through Repeated Interventions Over Time

**Benjamin M. Rottman (benjamin.rottman@yale.edu)**

Department of Psychology, Yale U., 2 Hillhouse Ave  
New Haven, CT 06520

**Frank C. Keil (frank.keil@yale.edu)**

Department of Psychology, Yale U., 2 Hillhouse Ave  
New Haven, CT 06520

## Abstract

How do we learn causal structures? All current approaches use scenarios in which trials are temporally independent; however, people often learn about scenarios unfolding over time. In such cases, people may assume that other variables don't change at the same instant as an intervention. In Experiment 1, participants were much more successful at learning causal structures when this assumption was upheld than violated. In Experiment 2, participants were less influenced by such temporal information when they believed the trials to be temporally independent, but still used the temporal strategy to some extent. People seem to be inclined to learn causal structures by connecting events over time.

**Keywords:** causal reasoning; causal structures; time

## Introduction

How do our concepts of event units influence how we learn causal structures? Despite the surge of research on causal structure learning, there has been little attention to how learners “connect” streams of information over time.

Existing theories of how people learn causal structures have focused on cases with events considered to be *independent*. For example, suppose we are trying to learn the causal relationships between three economic variables: employment, GDP, and consumption. Existing psychological theories suggest that one looks at the relationships among the variables across *many separate countries* to determine the causal structure. We call this strategy the *independent events* strategy because the countries are assumed to be independent.

An alternative approach is to pick one country and follow the three variables over time. We could track whether GDP goes up when employment goes up, etc. We call this strategy the *dependent events* or *temporal* strategy because the state of each variable is dependent on its prior state.

Psychologically, the temporal strategy may be pervasive and perhaps a default. As temporal beings we often perform or witness sequences of actions on one entity. For example, a car mechanic or computer technician can repair different components until the problem is solved. A psychotherapist can attempt to change one person's beliefs, emotions, and behaviors systematically over time. A physician can intervene on heart rate, breathing, and blood pressure to stabilize a patient. In many real-world situations we do interact with causal systems repeatedly over time, and thus

the temporal strategy may be common if not a default for learning causal structures.

In formal statistics we have developed specialized procedures for independent cases (e.g., between-subjects) and dependent cases (e.g., repeated-measures, time-series). Analogously, do people use different learning strategies for the two scenarios? In the following sections we detail the different inferences people might make.

## Interventions with Independent Trials

Consider first one prominent account of how people learn causal structures from interventions when trials are independent (e.g., Gopnik et al., 2004; Pearl, 2000; Steyvers et al., 2003). According to this model, when you intervene upon a variable such that you control its state, that variable is assumed to be independent from its other causes, but its effects are still dependent on that variable. Consider again the example of learning the causal relationships between employment ( $E$ ), GDP ( $G$ ), and consumption ( $C$ ). Pretend that a priori it is possible that any of these factors could influence or be influenced by any of the other factors. To learn the causal structure, one could intervene on each of the three variables to determine which other variables are influenced by (dependent upon) the intervention.

Suppose that the true causal structure is a chain;  $E$  influences  $G$ , which influences  $C$ ;  $E \rightarrow G \rightarrow C$ . If we could institute jobs-creation programs in 10 countries, we would expect them to have high  $G$  and  $C$ . If, hypothetically, we instituted a mass lay-off of government employees, we would expect comparatively low  $G$  and low  $C$ . These opposite interventions demonstrate how  $G$  and  $C$  are dependent on  $E$ . If we somehow selectively *boosted*  $G$  for 10 new countries, they would have high  $C$ , but the same  $E$  as if we *decreased*  $G$  for 10 other countries;  $C$  is dependent on  $G$  but  $E$  is not. And if we gave 10 countries a boost in  $C$ , and another 10 countries a decrease in  $C$ , the two countries should have the same  $E$  and  $G$ ; neither is dependent upon  $C$ .

If instead the true causal structure is a common cause such that  $E$  influences both  $G$  and  $C$ ,  $G \leftarrow E \rightarrow C$ , we would expect a different pattern of (in)dependence. If we increase or decrease  $G$ , the respective countries would have the same levels of  $E$  and  $C$  because they are independent of  $G$ .

This strategy can identify the precise causal structure because each causal structure has a different pattern of (in)dependence when the variables are intervened upon.

Importantly, however, this strategy requires that the observations be independent. This strategy does *not* look at whether one country's GDP *improves after* increasing employment compared to *before* (a within-subjects design). It only compares the outcome of countries with increased vs. decreased employment.

### Repeated Interventions Over Time

In contrast to the case just described, there are many scenarios in which a person intervenes repeatedly on one entity, and states of variables are fairly stable over time (e.g., car mechanic, physician). Consider a case in which we repeatedly intervene to increase or decrease  $E$ ,  $G$ , and  $C$  within the United States. Suppose that the true causal structure is  $E \rightarrow G \rightarrow C$ , and initially the country is in a recession and all three variables are low. If we start a job-creation program, we would expect  $G$ , and  $C$  to increase *compared to before the intervention*. Then, suppose that we decreased  $G$ . We would expect  $E$  to *stay high*, but  $C$  to decrease. Finally, suppose that we encouraged consumption. We would expect  $E$  and  $G$  to stay the same. In contrast, suppose that the true causal structure is  $G \leftarrow E \rightarrow C$ . Now, if we increase  $G$ , we would expect  $E$  and  $C$  to stay the same, but we would expect both to change if we intervened on  $E$ .

In sum, if we repeatedly intervene on one entity, we expect variables that are not influenced by the intervention to *remain constant*. If we intervene upon a variable  $X$ , and another variable  $Y$  changes *from the previous state*, it is a sign that  $X$  causes  $Y$ . If  $Y$  does not change when  $X$  is manipulated, it is a sign that  $X$  does not cause  $Y$ . These inferences are intuitive given the assumption that causes are generally stable and don't happen to change at the same moment that another cause is manipulated. This temporal assumption of "stability" is analogous to the atemporal assumption that interventions are independent of other causes (e.g., Pearl, 2000; see also Rottman & Ahn, 2009a).

### Testing Whether People Use the Two Strategies

The temporal strategy is very different from the strategy appropriate for independent observations. Only in the temporal case are the changes in variables over time important for learning causal structure and thus the order of the trials is critical.

To determine whether people are sensitive to the temporal information, we created pairs of data that have the same sets of 24 intervention trials, but with different trial orders. For example, consider the chain data in Figure 1. There are three variables ( $X$ ,  $Y$ , and  $Z$ ) and two possible values (0, and 1). Bold represents an intervention. For example, on Trial 1 for the useful chain condition,  $X$  was intervened upon and set to 1.  $Y$  and  $Z$  consequently have the value 1.

According to the independent trials strategy, both orders suggest the chain  $X \rightarrow Y \rightarrow Z$ . When  $X$  is intervened and set to 1,  $Y$  and  $Z$  are also 1. When  $Y$  is set at 1,  $Z$  is 1, but  $X$  can be either at 0 or 1 because  $X$  is not dependent on  $Y$ . Finally, if  $Z$  is set to 1,  $X$  and  $Y$  could both be 0 or 1 because they are independent of  $Z$ .

However, according to the temporal strategy, the two orders lead to very different inferences because the useful condition upholds the stability assumption but the misleading condition violates it. The "useful" condition suggests the  $X \rightarrow Y \rightarrow Z$  causal structure. Whenever  $X$  is changed,  $Y$  and  $Z$  also change (e.g. the transition from Trials 1 to 2). Whenever  $Y$  is changed,  $Z$  also changes, but  $X$  stays the same (e.g., Trials 2-3). When  $Z$  changes,  $X$  and  $Y$  stay the same (e.g., Trials 4-5). In contrast, misleading conditions were designed to suggest the presence of links that do not exist. For example, on Trial 2,  $Z$  is changed from 1 to 0, and  $X$  and  $Y$  also change to 0, suggesting that  $Z$  causes  $X$  and  $Y$ . Additionally, causal links are not consistent. On Trial 2,  $Z$  appears to cause  $X$  and  $Y$  to change to 0, but on Trial 3 it does not cause them to change back to 1. Finally, the existence of real links is obscured. For example, on Trial 5,  $X$  is changed from 0 to 1, but  $Y$  is already at 1, obscuring that  $X$  influences  $Y$ . In sum, the "misleading" condition suggests different links from the "useful" condition, and does not clearly identify one causal structure.

We used this order manipulation in two experiments. In Experiment 1, we tested whether people do in fact use the temporal strategy. In Experiment 2, we tested whether people appropriately switch between the two strategies based on the causal scenario.

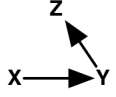
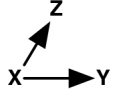
Trial	Chain 						Common Cause 							
	useful			misleading			useful			misleading				
	X	Y	Z	X	Y	Z	X	Y	Z	X	Y	Z		
1	<b>1</b>	1	1	1	1	1	1	1	1	1	1	1	1	1
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	<b>1</b>	1	0	0	1	1	1	1	1	1	1	1	1
4	0	0	0	0	<b>1</b>	1	1	0	1	0	0	0	0	0
5	0	0	<b>1</b>	1	1	1	1	1	1	1	1	1	1	1
6	0	0	0	0	0	0	0	1	1	0	0	0	0	0
7	<b>1</b>	1	1	1	<b>1</b>	1	1	1	1	1	0	0	0	1
8	1	0	0	0	0	0	0	0	0	0	1	1	0	0
9	1	<b>1</b>	1	0	0	1	0	1	0	0	0	0	0	1
10	1	1	0	0	<b>1</b>	1	0	0	0	0	1	1	0	0
11	1	1	<b>1</b>	1	1	1	0	0	1	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0	1	1	1
13	<b>1</b>	1	1	1	1	1	1	1	1	1	0	0	0	0
14	0	0	0	1	1	0	0	0	0	0	1	1	1	0
15	0	0	<b>1</b>	1	0	0	0	0	1	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0	0	1	1	1	0
17	0	<b>1</b>	1	1	<b>1</b>	1	0	1	0	1	0	1	0	1
18	0	0	0	0	0	0	0	0	0	0	0	1	0	0
19	<b>1</b>	1	1	1	1	1	1	1	1	1	1	0	1	0
20	1	0	0	1	1	0	1	1	0	0	1	0	0	0
21	1	<b>1</b>	1	1	0	0	1	1	1	1	1	1	1	1
22	1	1	0	0	0	0	0	1	0	1	0	0	0	0
23	1	1	<b>1</b>	1	1	1	1	1	1	1	1	1	1	1
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 1: Summary of Data for Two Causal Structures in Experiment 1.

## Experiment 1

In Experiment 1, we created a scenario in which one causal setup is repeatedly intervened upon over time. Thus participants would likely think that the temporal information was relevant. We presented participants with data generated by five causal structures. For each causal structure, there was a useful and misleading set of data. If participants use the temporal strategy, they will learn the causal structures more accurately in the useful condition.

### Methods

Twenty undergraduates completed the study for payment at \$10 per hour or partial course credit. Participants first read a cover story about three light bulbs. Participants were told that they would be instructed to turn on or off specific lights and should try to “learn how each light affects the others.”

Next, participants saw 10 scenarios created by crossing the Order of the Data (useful vs. misleading)  $\times$  Causal Structure (chain,  $X \rightarrow Y \rightarrow Z$ ; common cause,  $Y \leftarrow X \rightarrow Z$ ; common effect,  $X \rightarrow Z \leftarrow Y$ ; one link,  $X \rightarrow Y$ ,  $Z$  is unrelated; no links,  $X$ ,  $Y$ , and  $Z$ , are unrelated). The 10 scenarios were ordered in a Latin square grouped by causal structure such that each scenario appeared first for some participants.

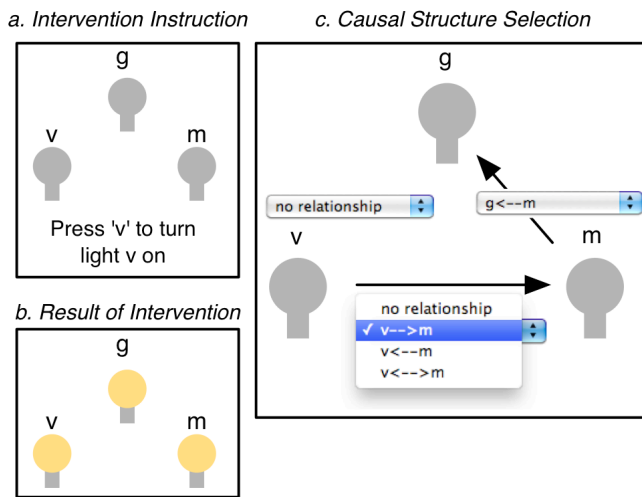


Figure 2. Example Screenshots from Experiment 1.

During each scenario, participants saw three light bulbs. Each bulb was named by a letter, and different letter triads were used across the 10 scenarios. Initially, all three bulbs were off. Then participants were instructed to intervene to turn on or off specific bulbs (e.g., Figure 2a). To intervene, participants pressed the key associated with the letter for the given bulb. After the intervention, participants observed the outcome of the intervention (which bulbs were on or off) for 2 seconds (e.g., Figure 2b). Then, while the bulbs were still visible, instructions appeared for the next intervention.

Each scenario had 24 interventions total, 8 per bulb; 4 on and 4 off. The data were determined in the following way. The causal relations were deterministic; when a bulb was intervened upon, all its effects (and all of their effects)

assumed the same value. Exogenous variables had a base-rate of .5. For the common effect structure, the effect was on if either of the causes was on.

For the “useful” conditions, the trials were ordered in a way that upheld the stability assumption explained in the introduction whereas the “misleading” conditions violated it. Figure 1 displays a summary of the data for the chain and common cause scenarios. The data for the other three causal structures can be obtained from the authors.

After each scenario, participants selected the causal structure that they believed to have generated the pattern of data for the given scenario (e.g., Figure 2c). Participants selected arrows indicating the direction of the causal relationships between the three light bulbs. For each pair of bulbs (e.g.,  $X$  and  $Y$ ), participants chose between “no relationship; neither light influences the other”, “ $X \rightarrow Y$ ;  $X$  influences  $Y$ ”, “ $X \leftarrow Y$ ;  $Y$  influences  $X$ ”, or “ $X \leftrightarrow Y$ ;  $X$  and  $Y$  both influence each other.” Participants did not receive feedback of the accuracy of their causal model. Finally, participants started the next scenario.

### Results

Accuracy in causal structure inferences was assessed in the following way. For each pair of bulbs,  $X$  and  $Y$ ,  $X$  can cause  $Y$  or not, and  $Y$  can cause  $X$  or not. Thus for each pair of bulbs, participants had the possibility of identifying zero, one, or two correct causal relations. Across the three bulbs in a given scenario, participants had the possibility of identifying zero to six correct causal relations.

For all of the five causal structures, participants identified more correct causal relations in the useful than misleading conditions  $t(19) > 8.32$ ,  $ps < .01$  (Figure 3), suggesting that they used the trial order for learning causal structures.

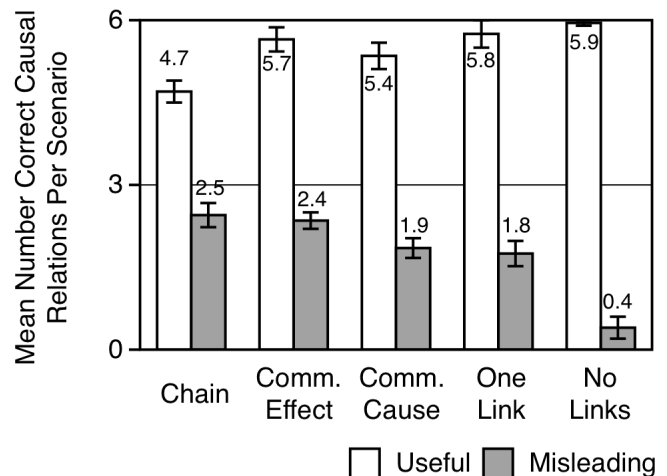


Figure 3: Mean Accuracy (Std. Errors) in Experiment 1.

There are two trends in participants’ mistakes. First, in the useful chain condition ( $X \rightarrow Y \rightarrow Z$ ), participants had difficulty learning that  $Y$  was a mediator between  $X$  and  $Z$ . This requires noticing that when  $Y$  is manipulated,  $X$  has no

influence on Z. Eighteen out of the 20 participants thought that X also caused Z directly, probably because when X was turned on and off, Z also changed state. Similar findings have been interpreted to suggest that people sequentially learn individual causal links rather than simultaneously learn an entire causal structure (Fernbach & Sloman, 2009).

Second, in the misleading conditions, participants frequently correctly identified true causal links, but they also mistakenly thought that other links existed. They often thought that links were bidirectional, even though they were just unidirectional. In the one link and no link conditions, they also frequently inferred relationships between variables with no causal relations. These inferences resulted in participants often misidentifying the majority of the causal links; the accuracy in all misleading conditions was below chance responding of 3, all  $t(19) > 2.4$ ,  $p < .03$ . However, these inferences make sense according to the temporal strategy; the misleading orders were designed so that variables that were not effects of a manipulated variable frequently change at the same time as the intervention, suggesting additional causal relationships.

In sum, the results strongly suggest that participants were sensitive to the order of the trials and were using the transitions between trials to infer causal relationships.

## Experiment 2

In Experiment 1, it was rational for participants to use a temporal strategy to learn causal structures because participants observed entities change over time. The purpose of Experiment 2 was to determine how flexibly people apply the temporal vs. independent strategies given different scenarios. We created two scenarios intended to give maximal cues to participants that the trials were either independent (analogous to a between subject design) or dependent (analogous to a within-subjects design). Previous studies have successfully used such a manipulation (Rottman & Ahn, 2009b). We then tested whether participants would infer different causal structures in useful vs. misleading orders. If participants use the temporal strategy for the dependent case, they would be more accurate in the useful than misleading order, as in Experiment 1. Additionally, if they do not use temporal information in the independent scenario, they would not have different levels of accuracy for the two orders.

## Methods

Sixteen students from the same population participated.

Participants first read a cover study story asking them to pretend that they are assistants in a biology lab studying hormones in amoebas. They would “produce” or “suppress” hormones by injecting chemicals into the amoebas and “learn how each hormone affects the others.” They were told that the “hormones work immediately... without any perceivable delay.”<sup>1</sup>

<sup>1</sup> This statement about no delay was intended to rule out the possibility of second order causal relationships (e.g., if Hormone A

Next, participants saw eight scenarios. Each scenario presented three hormones. “+” and “-” signs denoted the results of the hormones, presence and absence respectively. The eight scenarios were created by crossing Number of Amoebas (one vs. many) × Trial Order (useful vs. misleading) × Causal Structure (common cause,  $Y \leftarrow X \rightarrow Z$  vs. one link,  $X \rightarrow Y$ , Z is unrelated). The design was entirely within subjects. The 8 scenarios were ordered in a Latin square such that each scenario appeared first for some participants, and the scenarios were grouped by number of amoebas. The trial order and causal structure manipulations were the same as in Experiment 1, so the following paragraphs focus on the number of amoebas manipulation.

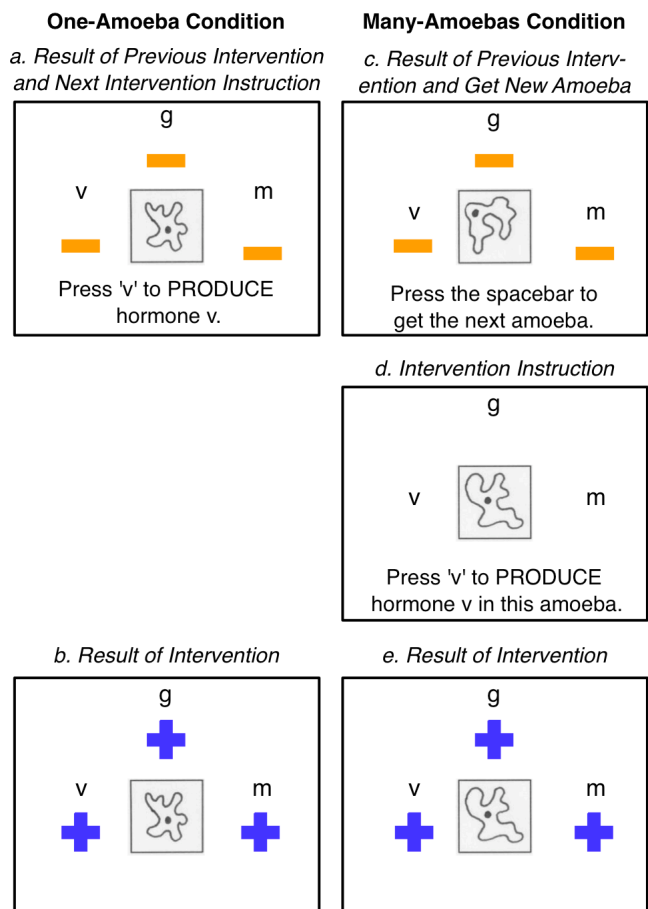


Figure 4: Example Screenshots from Experiment 2.

The one-amoeba condition, analogous to a within-subjects design, emphasized the dependent nature of the data. The one-amoeba procedures were similar to those in Experiment 1; participants repeatedly intervened on one amoeba. While the result of the previous intervention was displayed, participants were instructed to “PRODUCE” or “INHIBIT”

is produced and suppressed twice in a row, then Hormone B would be produced), which some participants reported in pretesting. In both the dependent and independent conditions, the interventions do work immediately after the intervention key is pressed.

a specific hormone (e.g., Press “y” to PRODUCE hormone y; e.g., Figure 4a). After the intervention, participants observed the result of the intervention for 2 seconds (e.g., Figure 4b). While the results were visible, instructions for the next intervention appeared. Additionally, a picture of one amoeba was present for the entire scenario to emphasize the repeated interventions on a single entity over time.

The many-amoebas condition, analogous to a between-subjects design, emphasized the independent nature of the data. Participants made 24 interventions on 24 different amoebas. After the results of a given intervention were displayed, participants were instructed to “Press the spacebar to get the next amoeba” (e.g., Figure 4c). When the spacebar was pressed, a picture of a new amoeba appeared. Simultaneously, the results of the intervention on the previous amoeba (“+” and “-” marks) disappeared (e.g. Figure 4d). We removed the previous results to make it perceptually difficult to track the changes of the hormones over time. Two seconds later, the prompt for the next intervention appeared (e.g., Press “y” to PRODUCE hormone y in this amoeba). When the intervention key was pressed, the hormone results appeared for the current amoeba (e.g., Figure 4e). All of these modifications were intended to signal that the hormones within one amoeba were independent of the hormones within other amoebas.

After each scenario, participants selected the causal structure that they believed to have generated the data.

## Results

The dependent variable was the same as in Experiment 1 – the number of correctly identified causal relations per scenario (zero to six).

A 2 (one vs. many amoebas)  $\times$  2 (trial order)  $\times$  2 (causal structure) repeated-measures ANOVA was performed. There was a main effect of trial order; participants correctly identified more causal relationships in the temporally useful than misleading orders,  $F(1,15)=45.28$ ,  $p<.01$ ,  $\eta_p^2=.75$  (Figure 5). However, the most critical result for this experiment is a significant interaction between number of amoebas and trial order,  $F(1,15)=12.61$ ,  $p<.01$ ,  $\eta_p^2=.46$ .<sup>2</sup> Though there was a large difference between the useful and misleading orders for the one-amoeba condition, there was a smaller difference between the many-amoebas conditions, suggesting that participants were less sensitive to the temporal order of trials in the many-amoebas condition. This finding makes sense if participants believed that the trials were independent in the many-amoebas condition.

However, even though participants used the temporal strategy less in the many-amoebas condition, they still used it to some extent; there was still a significant difference between the useful and misleading, many-amoebas conditions,  $t(15)=3.59$ ,  $p<.01$ . Furthermore, participants did

not simply transfer the temporal strategy from the one-amoeba condition; they were more accurate in the useful than misleading many-amoebas conditions even before experiencing the one-amoeba scenarios,  $t(7)=3.21$ ,  $p=.02$ .

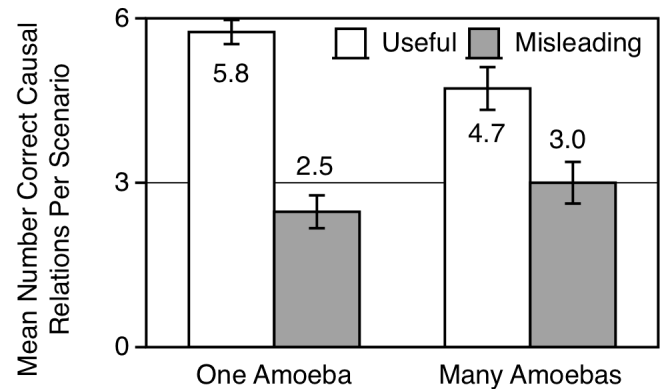


Figure 5: Mean Accuracy (Std. Errors) in Experiment 2.

There are two other important patterns. First, participants did worse in the many-amoeba than one-amoeba, useful condition,  $t(15)=2.57$ ,  $p=.02$ . This finding makes sense if participants were using the temporal strategy less in the many-amoebas condition. However, according to the independent trials strategies (e.g., Gopnik et al., 2004; Steyvers et al., 2003), participants should have been able to correctly identify the causal structures in the many-entity conditions. Second, participants were not even above chance in the many-amoebas, misleading condition,  $t(15)<1$ . Yet again, participants should have been able to identify the correct causal structures according to the independent trials strategy. The low accuracy in both many-amoebas conditions suggests that participants may have difficulty applying such statistical strategies.

In sum, participants are able to switch between the temporal vs. independent strategies to some extent based on knowledge of the learning scenario. However, even in the many-amoebas condition, participants used the temporal information to some extent, suggesting that it is a common strategy for learning causal structures.

## General Discussion

In two experiments, we demonstrated that people learn causal structures very well when entities are repeatedly manipulated over time (i.e. within-subjects or repeated measures situations). In Experiment 1, participants were much more accurate at learning causal structures when the data were ordered to reflect causes that are stable over time (don’t happen to change at the moment another variable is intervened upon), a plausible real-world assumption. In Experiment 2, participants were less sensitive to the temporal order of trials when they were given reason to believe that the trials were independent (i.e. between-subjects situation).

<sup>2</sup> The only other finding was a marginally significant interaction between causal structure and trial order,  $F(1,15)=4.03$ ,  $p<=.06$ ,  $\eta_p^2=.21$ . The difference between the useful and misleading orders was slightly larger for the common cause than one link conditions.

## Predominance of the Temporal Strategy

Why did participants in the many-amoebas condition in Experiment 2 still make use of the temporal information to some extent? There are two possible explanations. First, people may have still thought that the hormones within different amoebas were dependent upon one another. (For example, if all the amoebas were physically adjacent, perhaps hormones could mix across the amoebas.) Alternatively, people might have been able to learn that the trials were dependent from the data itself. In reality, in the many-amoebas, useful condition, the order was statistically dependent. For example, exogenous variables (e.g.,  $X$  in  $X \rightarrow Y \rightarrow Z$ ) only changed state when  $X$  was intervened upon. For long periods of time,  $X$  stayed the same (e.g., Trials 2-6 in Figure 1, Chain, Useful) even though its baserate is .5.

However, there is also a second possibility – the temporal strategy is likely simpler than the statistical strategies proposed for independent events (e.g., Gopnik et al., 2004; Steyvers et al., 2003). Thus, it is possible that people tend to use this strategy even in cases when the independent strategy is more appropriate. Perhaps the time-based strategy serves as a useful heuristic that is often accurate. In the real world, much of our causal reasoning involves manipulating and observing sequences of events unfolding over time (e.g., a car mechanic repairing different components until the problem is solved or a physician manipulating a patient's heart rate, breathing, and blood pressure to stabilize the patient). Given how frequently we engage in temporal reasoning, it may be hard to ignore temporal information such as the order of trials in these experiments even when we should for independent events.

## Learning Causal Structure from Temporal Delay

Lagnado and Sloman (2004, 2006; see also Burns & McCormack, 2009; Meder et al., 2008; White, 2006) showed how people use temporal *delays* when learning causal structures. For example, if you intervene upon  $X$ , and then  $Y$  appears, and later  $Z$  appears, you would likely infer  $X \rightarrow Y \rightarrow Z$ . This strategy pertains to the time course of how a causal signal propagates through a network and the order in which the reasoner becomes aware of the states of the nodes. This strategy is entirely consistent with the current one, and they likely often work in parallel in the real world. However, they are distinct. In the current studies, both of the non-manipulated variables appear simultaneously for all causal structures. Additionally, in the previous studies (e.g., Lagnado & Sloman, 2006), the trials were independent and were often randomized.

## Summary

Overall, people learn causal structures over time quite fluently and indeed seem biased to assume that this is the default mode of causal interpretation. Instead of treating trials as independent, which has been assumed by many approaches of causal structure learning, people weave together information across trials into larger event units.

The use of a temporal strategy can result in very quick and accurate causal structure learning when the trials are ordered in a temporally useful way. However, applying an incorrect causal strategy can result in substantially worse performance. For example, applying a more independent events strategy for events that were truly dependent and ordered in a useful fashion resulted in considerably worse performance than when participants applied the temporal strategy (Experiment 2). One intriguing possibility is that applying the temporal strategy when the events are truly independent could also likely result in reduced performance. Elaborating when and how people apply different learning strategies for diverse scenarios is an important future aim.

## Acknowledgments

This research was supported by an NSF Graduate Research Fellowship (Rottman) and NIH grant R37 HD023922 (Keil). The authors thank Samuel Tepper for programming, and Rachel Litwin and Brianna Sullivan for data collection.

## References

- Burns, P., & McCormack, T. (2009). Temporal information and children's and adults' causal inferences. *Thinking & Reasoning, 15*, 167-196.
- Fernback, P. M. & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 678-693.
- 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.
- Lagnado, D. A., & Sloman, S. A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory and Cognition, 30* (4), 856-876.
- 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.
- Meder, B., Hagmayer, Y., & Waldmann, M. R. (2008). Inferring interventional predictions from observational learning data. *Psychonomic Bulletin & Review, 15*, 75-80.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge Univ Pr.
- Rottman, B. M. & Ahn, W. (2009a). *Causal Inference when Observed and Unobserved Causes Interact*. Proceedings of the 31st Annual Conference of the Cognitive Science Society (pp. 1477-1482).
- Rottman, B. M. & Ahn, W. (2009b). *Causal Learning about Tolerance and Sensitization*. Psychonomic Bulletin and Review, 16, 1043-1049.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science, 27*, 453-489.