

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Causal Inference when Observed and Unobserved Causes Interact

Permalink

<https://escholarship.org/uc/item/7bf5p4g6>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN

1069-7977

Authors

Ahn, Woo-Kyoung
Rottman, Benjamin

Publication Date

2009

Peer reviewed

Causal Inference when Observed and Unobserved Causes Interact

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

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

Woo-kyoung Ahn (woo-kyoung.ahn@yale.edu)

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

Abstract

When a cause interacts with unobserved factors to produce an effect, the contingency between the observed cause and effect cannot be taken at face value to infer causality. Yet, it would be computationally intractable to consider all possible unobserved, interacting factors. Nonetheless, two experiments found that when an unobserved cause is assumed to be fairly stable over time, people can learn about such interactions and adjust their inferences about the causal efficacy of the observed cause. When they observed a period in which a cause and effect were associated followed by a period of the opposite association, rather than concluding a complete lack of causality, subjects inferred an unobserved, interacting cause. The interaction explains why the overall contingency between the cause and effect is low and allows people to still conclude that the cause is efficacious.

Keywords: Causal Learning; Causal Inference

Introduction

People often use co-variation between two observed events to infer causal relationships. However, when inferring the causal efficacy of an observed cause, assumptions and beliefs about unobserved causes are critical. For example, if an observed cause is confounded with an unobserved cause, normative theories (e.g., Cheng, 1997; Pearl, 2000) demand withholding inferences about the observed cause. Furthermore, if an unobserved cause produces an effect such that it is difficult for an observed cause to further increase the probability of the effect (a ceiling effect), a learner should adjust his/her calculation of the efficacy of the cause to account for the unobserved cause (Cheng, 1997).

Do people actually use beliefs about unobserved causes when making inferences about observed causes? Buehner and Cheng (1997; Buehner, Cheng & Clifford, 2003) found that people do adjust their causal efficacy estimates when unobserved causes produce floor or ceiling effects (but see Lober & Shanks, 2000; Vallee-Touragneau, et al., 1998, for alternative interpretations). Furthermore, recent studies have shown that people posit unobserved factors when an effect occurs in the absence of the observed cause (i.e., an unexplained effect; Luhmann & Ahn, 2007; Hagmayer & Waldmann, 2007). In summary, previous studies have shown that for floor/ceiling effects and unexplained events, peoples' inferences about unobserved causes influence learning about observed causes (but see Luhmann, 2005, for alternative possibilities).

The current study explores a third way that knowledge about unobserved causes can influence judgments of observed causes; observed and unobserved causes can interact to produce an effect. If a learner is unaware of such an interaction, he/she may incorrectly conclude that an observed variable has no causal relation with an effect. For example, consider a researcher studying if a new medicine reduces heart disease. Suppose that the medicine successfully reduced heart disease for half of the population, but the other half of the population had an undiscovered gene such that the medicine actually increased heart disease. In this scenario, even though the medicine is involved in a causal relationship with heart disease, there would be zero contingency between the two. This scenario poses a considerable challenge for causal learners; people cannot always look for interactions because there are simply too many possible interacting variables and there are always unobserved variables. When are people able to avoid overlooking real causal relationships that involve interactions with unobserved variables?

We propose that people may overcome this challenge if an unobserved variable is fairly stable over time and people assume it to be fairly stable. Consider an interaction scenario when there are two light switches connected to one light; the light is on if both switches are up or down, but the light is off if one switch is up and the other is down (the biconditional logical function). For this scenario, even if a person only knows about one of the switches, one may be able to infer that the switch influences the light. For example, suppose you enter a room for the first time and discover that when you flip a switch up, a light goes on, and when you flip it down, the light goes off (gray cells, Steps 1-4 in Figure 1a). If you assume that other potential causes of the light are fairly stable (and did not happen to change every time you flipped your switch), you would infer that the switch influences the light. Later (Steps 4-5), the light turns off without anyone touching the switch (perhaps your daughter flipped the other switch unknown to you; U in Figure 1a). Afterwards, when the switch is down, the light is on, and off when up (Steps 5-8). From this scenario, you might be very confident that your switch influences the light; there were two long periods when the status of the switch correlated with the status of the light. Additionally, because the light mysteriously turned off, you might infer an unobserved factor (I in Figure 1a) that interacts with your

switch, explaining the overall zero contingency between the switch and light.

However, inferring the observed switch to be efficacious may depend upon the order of the observations. For example, consider the same data from Figure 1a, rearranged as in Figure 1b, which could result if an unobserved factor changed frequently. Initially, the switch is down and the light is off (Step 1). In Step 2 the switch is flipped up, but the light still stays off. In order to believe that the switch is causally efficacious, one must infer that at the moment the switch was flipped, an unobserved factor coincidentally changed and counteracted the effect of the observed switch, as specified under column “U” (unobserved interacting factor). Then, in Step 3, the light turns on without flipping the switch, and so on. Thus, for the situation shown in Figure 2, it would be extremely difficult to infer the switch to be causally efficacious; the switch cannot be the sole cause of the light because there is zero contingency with the light. Furthermore, it would be difficult to infer it as part of an interaction because doing so would require inferring an unobserved factor operating as specified under row “U,” which is counterintuitive; the unobserved interacting factor is highly unstable and (intuitively) exceedingly complicated to track. Instead, it seems likely that people would infer an unobserved factor that is *entirely* responsible for the light. Such a factor would be perfectly correlated with the light, as specified under column “I” (inferred factor).

a. Grouped Scenario

Steps	1	2	3	4	5	6	7	8
Switch	Dn	Up	Dn	Up	Up	Dn	Up	Dn
Light	Off	On	Off	On	Off	On	Off	On
U	Up	Up	Up	Up	Dn	Dn	Dn	Dn
I	Up	Up	Up	Up	Dn	Dn	Dn	Dn

b. Ungrouped Scenario

Steps	1	2	3	4	5	6	7	8
Switch	Dn	Up	Up	Dn	Dn	Up	Up	Dn
Light	Off	Off	On	On	Off	Off	On	On
U	Up	Dn	Up	Dn	Up	Dn	Up	Dn
I	Off	Off	On	On	Off	Off	On	On

Figure 1: Double Light Switch

Note: “U” is an unobserved, interacting switch and “I” is a factor learners are likely to infer. Gray is when U is “Up,” highlighting the (in)stability of U. “Dn” means “down.”

These two examples were meant to demonstrate that if observations are grouped (reflecting stable background causes, e.g., Figure 1a), rather than inter-mixed (reflecting unstable background causes, e.g., Figure 1b), people may be more likely to infer that an interaction is taking place and that the observed cause is efficacious rather than that an unobserved factor is entirely responsible for the effect. Such findings would suggest not only that people make sophisticated inferences about interactions, but also that people tend to assume that unobserved causes are stable.

In two experiments, we manipulated the grouping of observations. In the grouped condition, the trials supporting an association between one state of the cause and effect (gray cells in Figures 1 and 2) are grouped together, and those supporting the opposite association (white cells in Figures 1 and 2) are grouped together. In the ungrouped condition, these two types of observations are inter-mixed. Although overall contingency is identical between the two conditions, participants may be more likely to infer an interacting unobserved factor and an efficacious observed cause in the grouped rather than ungrouped condition. We examined whether grouping of observations would result in higher causal efficacy ratings (Experiment 1) and higher ratings that an observed cause interacts with an unobserved factor (Experiment 2).

Experiment 1

Experiment 1 compares the grouped and ungrouped conditions across three levels of contingency. We predicted that people would infer stronger causal efficacy of observed causes in the grouped than in the ungrouped condition.

Methods

Thirty-six undergraduates from Yale University completed the study for payment at \$10 per hour. Participants first read the following cover story:

On the following screens you will see machines with a lever that can be set to two positions and you will see the toys that the machines produce (e.g. square or triangle blocks). The position of the lever and the shape of the blocks change over time. I would like you to determine whether the position of the lever affects the shape of the blocks. Though the following scenarios may look similar, please pay attention because there are differences. Also, please note that if the lever affects the blocks, it happens immediately.

Next, participants saw six scenarios created by crossing Grouping (grouped vs. ungrouped) × Contingency ($\Delta P = .25, .5, \text{ or } .75$).¹ The six scenarios were ordered in a Latin square design (grouped by contingency) such that each of the six scenarios appeared first for some participants.

During each scenario, participants viewed a video of a lever changing (between left or right) and blocks changing between two shapes (e.g., square or triangle). Each of these binary values will be denoted as 0 and 1 henceforth (e.g., 00 trial means a lever set to the left and the shape being square). Each scenario had 16 trials, each of which appeared for 2 seconds followed immediately by the next for each of the three contingencies. C and E represent the cause (lever) and effect (shape of block), respectively.

¹ ΔP is the probability of an effect when a cause is present minus the probability of the effect when the cause is absent (Jenkins & Ward, 1965). For these stimuli, since the lever and shape of blocks do not have presence/absence states, ΔP was defined as $p(e=1|c=1) - p(e=1|c=0)$. For the data summarized in Figure 2, the absolute value of this definition is unchanged if 0/1 are swapped.

Experiment 1

1 ————— Trial Number —————> 16

$\Delta P = .25$: Grouped

C	1	0	1	0	1	0	0	1	0	1	0	1	0	1	0	1
E	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
U	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1

$\Delta P = .25$: Ungrouped

C	0	1	1	0	1	0	1	0	1	0	1	0	0	1		
E	0	0	1	1	1	0	0	0	1	1	1	0	0	0	1	1
U	1	0	1	0	1	1	0	1	1	0	1	1	0	1	0	1

$\Delta P = .50$: Grouped

C	1	0	1	0	0	1	0	1	0	1	0	1	0	1	0	1
E	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
U	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1

$\Delta P = .50$: Ungrouped

C	1	0	0	1	0	1	1	0	1	0	0	1	0	1	0	1
E	1	1	0	1	0	0	1	0	1	1	0	1	0	0	0	1
U	1	0	1	1	1	0	1	1	1	0	1	1	1	0	1	1

$\Delta P = .75$: Grouped

C	1	0	0	1	0	1	0	1	0	1	0	1	0	1	0	1
E	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
U	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1

$\Delta P = .75$: Ungrouped

C	1	0	1	0	1	0	1	0	1	0	0	1	0	1	0	1
E	1	0	1	0	0	0	1	0	1	0	1	1	0	1	0	1
U	1	1	1	1	0	1	1	1	1	1	0	1	1	1	1	1

Experiment 2

1 ————— Trial Number —————> 20

$\Delta P = 0$: Grouped

C	1	0	1	0	1	1	0	1	0	0	1	0	0	1	0	1	1	0	0	1
E	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
U	1	1	1	1	1	0	0	0	0	1	1	1	0	0	0	0	1	1	0	0

$\Delta P = 0$: Ungrouped

C	0	0	1	1	0	1	0	0	0	1	0	1	1	1	1	0	1	0	1	0
E	0	1	0	1	1	1	1	0	1	1	0	0	0	0	1	1	1	0	0	0
U	1	0	0	1	0	1	0	1	1	0	0	0	1	0	1	1	0	1	0	1

Figure 2: Summary of Stimuli in Experiments 1 and 2. Note: “U” represents an unobserved, interacting factor not shown to participants. Cells are gray when $U=0$.

In the grouped conditions, 01 and 10 trials (i.e., gray cells in Figure 2) appeared in one cluster, (e.g., trials 1-6 for $\Delta P = .25$), and 11 and 00 trials (i.e., white cells in Figure 2) appeared in another cluster. In the ungrouped conditions, the four types of trials were intermixed. In Figure 2, the “U” column shows what the value of an unobserved factor would need to be in order to postulate that the observed cause and unobserved factor interact to produce the effect. As illustrated here, in the ungrouped condition, one must infer a highly unstable unobserved factor to infer an interaction, which would be very difficult for participants.

Because people often base causal efficacy ratings more on initial than final trials (e.g., Dennis & Ahn, 2001), the trials were presented in the reverse order for half the participants. After each scenario, participants answered one causal efficacy question, “To what extent does the lever affect the shape of the blocks?” on a sliding scale from “The lever **did not** affect the shape at all” to “The lever **strongly** affected the shape of the blocks,” later recoded to 0-100 for analysis.

Results

Participants’ average causal efficacy ratings for the six scenarios are presented in Figure 3. The pattern of results is consistent regardless of the order of the six scenarios and regardless of whether the order of trials within a scenario was reversed; all analyses collapse across these factors.

As can be seen in Figure 3, the most dramatic finding is that participants gave higher causal efficacy ratings for the grouped than ungrouped conditions. This is presumably because, in the grouped conditions, participants inferred the lever to influence the shape of the block through an interaction with an unobserved variable. In a 2 (grouping) \times 3 (contingency) repeated-measures ANOVA, the main effect of grouping was significant, $F(1,35)=75.90$, $p < .01$, $\eta_p^2 = .69$. Furthermore, the main effect of contingency was significant, $F(2,70)=7.23$, $p < .01$, $\eta_p^2 = .17$, but there was no significant interaction, $F(2,70) < 1$.

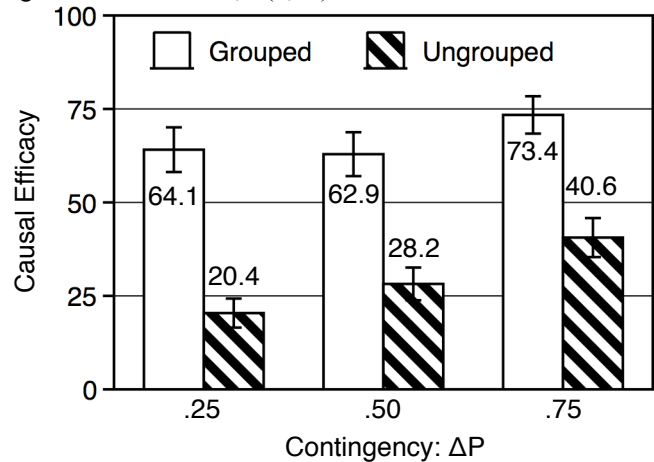


Figure 3: Causal Efficacy Ratings and Std. Errors in Exp. 1.

Follow-up tests reveal that for each of the three contingencies, participants gave higher causal efficacy ratings for the grouped than ungrouped conditions, all t 's(35) > 5.48 , all p 's $< .01$. To determine the effect of contingency, the same ANOVA as above was run testing for a linear effect of contingency. A significant linear effect was found, $F(1,35)=13.59$, $p < .01$, $\eta_p^2 = .28$; participants gave higher causal efficacy ratings for higher contingency levels. There was no interaction between the linear effect of contingency and grouping, $F(1,35)=1.33$, $p = .26$, $\eta_p^2 = .04$. In summary, these results suggest that if observations are ordered such that reasoning about an interacting,

unobserved factor is more feasible, learners judge the causal efficacy of the observed cause to be greater than if not.

Yet, a possible alternative explanation for the current findings is that the grouping may have allowed participants to focus only on one set of data (e.g., 10 and 01 trials) and ignore the contradicting data (e.g., 11 and 00 trials). If so, participants would have inferred that the observed cause influenced the effect alone, rather than in combination with an unobserved cause. Since we asked questions only about the observed causes, as in almost all previous causal learning experiments, we cannot tell whether the current results were obtained due to the participants' inferences about an interaction with unobserved causes. The next experiment addresses this issue.

Experiment 2

In Experiment 2, we explicitly asked participants about their belief in an interaction. In this way, Experiment 2 attempts to demonstrate more explicitly that the stronger causal efficacy judgments found in Experiment 1 were due to participants' inferring an interaction with an unobserved cause. As explained in the Introduction, when observations are grouped so as to encourage participants to initially believe that the observed cause is causally responsible for the effect, they would likely attribute the subsequent contrary evidence to an interaction with an unobserved factor. Even when the overall contingency is zero (Experiment 2, grouped condition), if participants infer an interaction between the observed and unobserved causes, they may still judge the observed cause to influence the effect; it influences the effect *in combination with* the unobserved factor, not alone. However, when participants cannot easily infer that the observed and unobserved causes interact (ungrouped conditions), they should judge the observed cause to be less efficacious; after all, there is no contingency. Also, importantly, participants in both conditions should understand that the observed cause does not influence the effect *alone*. In the grouped condition, it influences the effect in combination, and in the ungrouped condition, it does not influence the effect at all.

Methods

There were twenty-nine participants from the same population as Experiment 1.

Experiment 2 only had two conditions, grouped and ungrouped. In the grouped condition, three groups of 00 and 11 trials and three groups of 01 and 10 trials alternated, while in the ungrouped condition, the four types of trials were more intermixed, as shown in Figure 3. The amount of grouping in the grouped condition is much less than in Experiment 1, offering a more rigorous test of participants' ability to infer an interaction. The two conditions comprised identical data with zero contingency, and both were presented to all participants in a counterbalanced order.

The presentation of the scenarios was the same as in Experiment 1 except for the following changes. The number of trials was increased to 20. In order to avoid any potential

ambiguities about the trial order, each scenario also had a header present for the entire scenario stating, "Below you will see **one machine** tested for **twenty consecutive trials.**" A picture of a machine was also present for all 20 trials, and a different picture was used for the two conditions.

After each scenario, to understand the relationship between their beliefs about the interaction and the main effects of the observed cause, we had participants rate their agreement with three causal efficacy statements from 1 ("Absolutely Disagree") to 9 ("Absolutely Agree"). The statements were:

- 1) "The lever **alone** influenced the shape of the blocks."
- 2) "A combination of the lever and some other factor influenced the shape of the blocks."
- 3) "The lever had **no influence** on the shape of the blocks."²³

Results

Participants thought that the lever influenced the shape of the blocks in combination with an unobserved factor more in the grouped than ungrouped scenario, $t(28)=4.20$, $p<.01$ (left panel of Figure 5). Restated, grouping increased participants' inferences of an interaction.

When participants did not believe in an interaction (i.e., ungrouped condition), they seemed to use the zero contingency as evidence that the lever did not influence the effect. Thus, participants gave significantly higher ratings that the cause did not influence the effect in the ungrouped compared to grouped conditions, $t(28)=3.65$, $p<.01$ (middle panel of Figure 4). In fact, in both the grouped, $r(29)=-.51$, $p<.01$, and ungrouped conditions, $r(29)=-.50$, $p<.01$, there were significant negative correlations between judgments of an interaction and judgments that the lever had no influence.

Participants also understood that the lever did not influence shape alone (right panel of Figure 4), and there was no difference between conditions, $t(29)=.88$, $p=.39$. These findings are consistent with our accounts: In the

² These three questions essentially ask participants to choose between the three causal models ($C \rightarrow E$, $C \rightarrow E \leftarrow U$, and $U \rightarrow E$; C=cause [lever], E=effect [shape], U=unobserved factor). Note that no two of these questions are exact opposites of one another. For example, if a person believes that C alone influences E (agreement with Question 1), he/she should disagree that C and U combine to produce E (disagreement with Question 2). However, if a person believes that C does not influence E (agreement with Question 3), he/she should disagree with both Questions 1 and 2. These questions are designed to allow participants to show which of the three options they agree more with, and participants may potentially be agnostic across the three.

³ Afterwards, participants predicted the shape of the block given that the lever was set to the left/right. These questions tested hypotheses that are not the main focus of the current report and are compared with between-subject conditions not reported here. Thus, they will not be further discussed. We do not think they had any effect on the current results as the findings are consistent with a between-subjects analysis of questions answered prior to the unreported questions. Thus, they will not be further discussed.

grouped condition, participants understood that the lever influenced shape in combination with an unobserved factor, not alone. In the ungrouped condition, participants believed that the lever had little influence in combination or alone. Because grouping had no effect on judgments of whether the lever influenced shape alone, it seems that inferences about the unobserved factor rather than grouping per se were responsible for the results in Experiment 1.

To summarize, although participants observed zero contingency, they could reason that the observed cause was causally efficacious by way of interacting with another factor that was not even observed. Such sophisticated inferences were more common in the grouped than in the ungrouped condition.

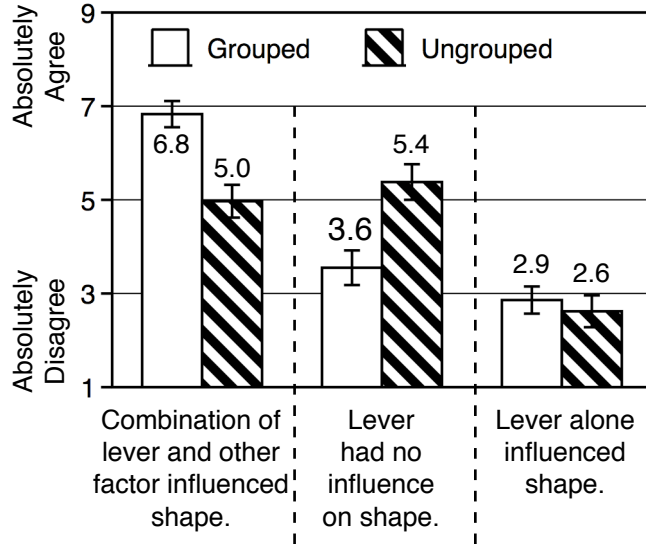


Figure 4: Causal Efficacy Ratings and Std. Errors in Exp. 2.

General Discussion

In two experiments, we demonstrated that when data are grouped reflecting a stable unobserved cause, people could infer that an observed cause interacted with an unobserved factor to produce an effect. They further rated observed causes that they believed to participate in an interaction as more efficacious than causes that they did not believe participate in an interaction but had the same contingency.

How did participants infer an unobserved interacting cause? In the grouped conditions, there was an initial group of observations that would allow people to suspect that the observed cause was causally efficacious (e.g., Steps 1-4 in Figure 1a). Then, when the direction of the relationship changed (e.g., Steps 5-8 in Figure 1a) participants may interpret this change as evidence of an interaction with an unobserved cause.

Is inferring such an interaction rational? On one hand, inferring an unobserved cause may be an irrational form of motivated reasoning; that is, in order to perpetuate an initial hypothesis (e.g., the switch influences the light) in the face of contrary evidence, the learner concocts an interacting cause to dismiss that evidence. On the other hand, such an inference is based on the assumption that unobserved

background causes are stable and do not change erratically (e.g., the status of an unobserved cause is the same through Steps 1-4, and Steps 5-8 in Figure 1a), which may be a reasonable and rational assumption to make in the real world. Indeed, when previous researchers have discussed the value of intervening upon a cause (as opposed to passively observing the changes of the cause) for causal learning, their arguments have relied upon a similar assumption that interventions are not confounded with changes to other variables (Pearl, 2000; Woodward, 2003). If this assumption reasonably approximates our environment, the kind of inferences shown in the current study may guide learners toward accurate answers.

Why did participants provide such high causal efficacy ratings for the grouped conditions? For example, in a grouped condition in Experiment 1, the average causal efficacy rating was 64 out of 100 even though ΔP was only .25. When judging causal efficacy, participants may have tried to use periods of stability of the unobserved cause to calculate the causal efficacy of the observed cause holding the unobserved factor constant. For example, consider the $\Delta P=.5$, grouped condition from Experiment 1, summarized in Figure 2. For the first four trials, the shape of the block changes when the lever changes. Then, from Trials 5-16, the shape also changes when the lever changes. Thus, the conclusion within both of these periods when the unobserved cause was likely inferred to be constant is that the lever influences shape. However, in the ungrouped condition, our participants likely believed that the unobserved cause was not constant for very long. Consequently, they might have used the overall contingency between the cause and effect to estimate causal efficacy. Previous studies have shown that when learning causal relations, people can “condition” on an *observed*, alternative cause (e.g., Spellman, 1996). In the current study, it appears as if people are simultaneously conditioning on both states of the *unobserved* cause. Another way to explain this phenomenon is that participants may have used the transitions between trials to infer causal efficacy. In the grouped conditions, there were many transitions during which the cause and effect both changed state (e.g., 11 to 00 or 10 to 01). However, in the ungrouped conditions, there were many trials when the effect changed state without the cause (e.g., 00 to 01), suggesting that an unobserved factor changed. There are also many transitions when a change in the cause failed to produce a change in the effect (e.g., 00 to 10), suggesting that the cause does not influence the effect. We intend to investigate such inferences in future studies.

Implications for Models of Causal Learning

Existing models of causal learning fail to capture our participants’ inferences for various reasons. First, the current study shows that people are sensitive to the temporal order of events, yet previous studies have largely neglected temporal order. Even in studies that present trials sequentially, each trial typically presents a separate case (e.g., a person taking/not taking medicine and

developing/not developing heart disease). In the current study, one machine with one lever is repeatedly tested across a period of time. Consequently, participants can make rich inferences about transitions between trials (explained above) unavailable in the previous experiments.

Yet, existing models cannot account for the role of temporal information in causal learning. Some models (e.g., Jenkins & Ward, 1965) fail because they aggregate over all trials regardless of order. Consequently, they do not differentiate between the grouped and ungrouped conditions. Models that continually update their causal efficacy estimate after each trial (e.g., Luhmann & Ahn, 2007; Rescorla & Wagner, 1972) fail in a different respect. When the cause produces the effect, these models would calculate a positive causal efficacy, and when the cause inhibits the effect, they would calculate a negative causal efficacy. Consequently, depending on the stability of the unobserved cause (how frequently it changes between producing vs. inhibiting the effect), the causal efficacy judgments would cycle back and forth indefinitely. Future models may need to focus on how causes change over time.

Second, existing models also fail to capture inferences about the interaction between the observed and unobserved causes. For instance, the power PC theory (Cheng, 1997) requires that unobserved causes do not interact with observed causes. Previous attempts to handle interactions between two *observed* causes and an effect (e.g., Cheng & Novick, 1992; Novick & Cheng, 2004) also cannot account for the current findings. Unless a learner has an a priori reason for believing that an interaction is taking place, he or she must first infer whether there is an interaction with an unobserved variable. Yet none of the existing models are able to detect interactions with unobserved causes.

Third, the methods used to assess causal efficacy in the current study offer another significant implication for models of causal learning. Traditionally, causal efficacy has been assessed and modeled in a directional manner (generative or inhibitory causal relationship; e.g., “to what extent does X *cause* Y” or “to what extent does X *inhibit* Y”). However, the phenomenon demonstrated in the current study would not have been captured by questions specifying the direction. For instance, in the double light switch scenario in Figure 1a, sometimes the switch causes the effect and sometimes it inhibits the effect. If assessed in the traditional way, a learner would likely say that the observed variable does not *cause* the effect. Instead, the current study measured the overall causal influence of X on Y (regardless of direction), which is more general and potentially more sensitive to diverse types of causal relationships such as interactions. Most models, however, compute directional causal efficacy, and therefore, cannot capture the type of causal learning demonstrated in the current study.

Acknowledgments

This research was supported by an NSF Grad. Research Fellowship (Rottman) and NIMH Grant R01 MH57737 (Ahn). The authors thank Rachel Litwin for data collection.

References

- Buehner, M. J., & Cheng, P. W. (1997). Causal induction: The Power PC Theory versus the Rescorla-Wagner Model. In the *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology*, *29*, 1119-1140.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.
- Hagmayer, Y., & Waldmann, M. R. (2007). Inferences about unobserved causes in human contingency learning. *Quarterly Journal of Experimental Psychology*, *60*, 330-355.
- Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, *79*, 1-17.
- Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, *107*, 195-212.
- Luhmann, C. C. (2005). Confounded: Causal inference and the requirement of independence. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1355-1360). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.
- Luhmann, C. C., & Ahn, W. (2007). BUCKLE: A model of unobserved cause learning. *Psychological Review*, *114*(3), 657-677.
- Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, *111*, 455-485.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.) *Classical conditioning II: Current research and theory*. New York, NY: Appleton-Century-Crofts.
- Spellman, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, *7*, 337-342.
- Vallee Tourangeau, F., Murphy, R. A., Drew, S., & Baker, A. G. (1998). Judging the importance of constant and variable candidate causes: A test of the power pC theory. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *1*, 65-84.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.