

Preventative Scope in Causation

Christopher D. Carroll (cdcarroll@ucla.edu)

Department of Psychology, UCLA

Patricia W. Cheng (cheng@lifesci.ucla.edu)

Department of Psychology, UCLA

Abstract

Some preventers only stop an effect when it is being produced by certain causes. For example, nasal spray prevents headaches caused by a cold but not headaches caused by dehydration or stress. Thus, preventers differ in preventative scope: the range of circumstances across which a preventer operates. An experiment indicated that people are sensitive to differences in preventative scope and that participants are more likely to generalize prevention when the preventer has a broad preventative scope. Additional evidence suggested that people take preventative scope into account when attempting to explain how prevention operates.

Keywords: causality; causal power; prevention; hidden causes

Preventative Scope

What does it mean for one thing to prevent another? While many models of causal reasoning include prevention (e.g., Cheng, 1997; Griffiths & Tenenbaum, 2005), few studies have examined how well prevention in these models corresponds to everyday notions of prevention. One aspect of prevention that is found in everyday reasoning but poorly represented in formal models is preventative scope: the range of circumstances across which the preventer works.

We define preventative scope relative to the causes of the effect. A preventer with a *broad preventative scope* stops the effect no matter what the cause. A preventer with a *narrow preventative scope* stops the effect only when certain causes are responsible for producing the effect. The difference between broad prevention and narrow prevention can be illustrated by contrasting the influences of aspirin and nasal spray on headaches. Aspirin prevents headaches caused by colds, headaches caused by dehydration, and headaches caused by stress. Nasal spray prevents headaches caused by colds, but not headaches caused by dehydration or stress. Thus, aspirin has a broad preventative scope and nasal spray has a narrow preventative scope.

An experiment was conducted to investigate preventative scope. The goals of the experiment were to determine (1) whether people are sensitive to preventative scope, (2) how preventative scope is generalized to situations where there is a novel generative cause, and (3) how people interpret or explain preventative scope.

Broad and narrow prevention

In the current investigation, we restrict ourselves to the extremes of preventative scope: we contrast a narrow preventer whose preventative scope includes only a single

generative cause and a broad preventer whose preventative scope includes all of the generative causes of the effect.

Preventative scope can be inferred by observing the frequency of the effect as a function of its causes and the preventer. Even in the simplest situation where only one generative cause is known, a broad preventer and a narrow preventer predict different patterns of covariation between the generative cause, the preventer, and the effect. To see why, consider the influence of the preventer when the generative cause is absent. When an effect occurs in the absence of a known cause, its occurrence is often attributed to one or more unknown or unobserved causes (Hagmayer & Waldmann, 2007; Luhmann & Ahn, 2007; Saxe, Tenenbaum, & Carey, 2005). By definition, a broad preventer stops these unknown causes from producing the effect, and a narrow preventer does not. Broad prevention should influence the frequency of the effect even when the generative cause is absent, whereas narrow prevention should not do so.

More formally, for causes and effects that are binary (either present or absent), it is possible to derive the probability of the effect as a function of the generative cause, the preventer, and the preventative scope of the preventer. We do so by adopting the assumptions in Cheng (1997) and adding the assumption of preventative scope. Letting p_c be the causal power of the generative cause, p_p be the causal power of the preventer, and p_a be the causal power of the causal background (representing causal power and frequency of the unknown causes of the effect in the current context), we obtain the equations for broad prevention and narrow prevention that are shown in Table 1.

Table 1: Probability of effect as a function of cause c and preventer p

	broad prevention	narrow prevention
$P(e \neg c, \neg p)$	p_a	p_a
$P(e c, \neg p)$	$p_a + p_c - p_a p_c$	$p_a + p_c - p_a p_c$
$P(e \neg c, p)$	$p_a (1 - p_p)$	p_a
$P(e c, p)$	$p_a (1 - p_p) + p_c (1 - p_p) - p_a p_c (1 - p_p)$	$p_a + p_c (1 - p_p) - p_a p_c (1 - p_p)$

The differences between the formulas for broad and narrow prevention arise because only a broad preventer reduces the probability of the effect when the effect is produced by an unknown cause in the causal background.

As long as the effect occurs in the absence of the generative cause (i.e., $pa > 0$), narrow prevention and broad prevention make different predictions. For a narrow preventer, the preventer does not influence the probability of the effect when the generative cause is absent $[P(e|\sim c, \sim p) = P(e|\sim c, p)]$. When the preventer is broad, however, the preventer reduces the probability of the effect even when the generative cause is absent $[P(e|\sim c, \sim p) > P(e|\sim c, p)]$.

One goal of the current investigation is to test whether people are sensitive to the differences between broad prevention and narrow prevention. Do people notice these differences, and are these differences used to make inferences about preventative scope? In the current experiment, we showed participants some data and manipulated whether it was consistent with broad or narrow prevention. Then we tested inferences about preventative scope by asking whether the preventer would stop the effect when it is caused by a novel generative cause.

Interpreting broad and narrow prevention

The differences between broad and narrow prevention beg an explanation: why do some preventers stop the effect regardless of the cause while others do not? Narrow prevention seems to be especially problematic. How is it possible for a narrow preventer to stop one cause from producing the effect but not others?

One explanation proposes that the generative cause and the narrow preventer both influence the effect through a common mediating variable. That is, the generative cause causes a mediating variable that in turn causes the effect. The narrow preventer also acts upon the mediating variable rather than acting upon the effect directly. This sort of




Figure 1: A potential explanation of narrow prevention. White arrows indicate generative causation. Black arrows indicate prevention.

explanation underlies the narrow preventative scope of nasal spray (see Figure 1). A cold causes sinus inflammation and sinus inflammation causes headaches. Nasal spray reduces sinus inflammation, and thus prevents colds from producing headaches. Other causes (e.g., dehydration) produce headaches through other means and are unaffected by nasal spray. This explanation sits well with the intuition that a narrow preventer interrupts or blocks the mechanism (instantiated as a mediating variable) through which the cause produces the effect.

On the other hand, broad prevention may suggest another causal explanation: a broad preventer may act directly upon the effect. Existing models of causal reasoning can represent broad prevention in this manner (e.g., Cheng, 1997).

If people are sensitive to preventative scope, then inferences about preventative scope may determine which of these explanations someone prefers. The analysis above suggests that people shown narrow prevention will be more likely to endorse an explanation that involves mediation.

Method

Participants were given a cover story and then presented with observations. We manipulated whether the observations were consistent with broad or narrow prevention. Participants were then asked a series of questions intended to reveal the circumstances under which they expected the preventer to be effective and how they interpreted the prevention.

Participants

Forty undergraduates at the University of California, Los Angeles (UCLA) participated to obtain course credit in a psychology course.

Materials

Data generated for the narrow prevention condition and the broad prevention condition are shown in Table 2. There were fifty observations for each possible combination of the generative cause and preventer. The broad prevention data were generated by setting $pc=.75$, $pp=.5$, and $pa=.1$ and applying the broad prevention equations. The narrow prevention data were generated by setting $pc=.75$, $pp=.625$, and $pa=.1$ and applying the narrow prevention equations.

The narrow preventer had a higher causal power (.625) than the broad preventer (.5) in order to control for the overall efficacy of the preventer. That is, these causal powers were chosen so that the narrow preventer and broad preventer were equally effective after collapsing over the presence and absence of the generative cause. In each condition, the effect occurred 50 / 100 times when the preventer was absent and 25 / 100 times when the preventer was present. This control was necessary to exclude some alternative explanations for the predicted results. For example, the overall efficacy of the preventer might be used as a heuristic to infer the number of variables mediating a

causal relationship. If so, then participants may prefer explanations with mediation when the overall efficacy of the preventer is low. Similarly, the generalization of prevention to novel causes might depend on the overall efficacy of the preventer. By controlling for the overall efficacy of the preventer, we were able to exclude these explanations and isolate the influence of preventative scope.

Table 2: Frequency of the effect as a function of the generative cause c , preventer p , and experimental condition

	broad prevention	narrow prevention
$\sim c, \sim p$	10 out of 50	10 out of 50
$c, \sim p$	40 out of 50	40 out of 50
$\sim c, p$	5 out of 50	10 out of 50
c, p	20 out of 50	15 out of 50

There were two cover stories that asked the participants to imagine themselves as researchers at a medical company. Each cover story first introduced an effect. Participants were then told that they would research how two candidate variables influenced the occurrence of the effect. A colleague at the research company was said to have investigated one of these candidate variables (the *generative cause*) and found that it produced the effect indirectly via another variable (the *mediating variable*). That is, the generative cause caused the mediating variable, and the mediating variable caused the effect. The mediating variable was described as rare and poorly understood. No information was given about the other candidate variable (the *preventer*), so participants did not know whether or how it influenced the effect prior to viewing the data. Finally, each cover story was associated with a *novel cause* that produced the effect. The novel cause was not mentioned in the cover story, but was used later in the experiment to test the generalization of the preventer. One of the cover stories is provided below:

Imagine that you work for a drug company that develops headache medications. The company has asked you to investigate pane fruit and asmine juice. In the rainforest, people eat pane fruit and drink asmine juice because they are quite delicious. However, you are more interested in understanding the effects of eating pane fruit and drinking asmine juice on headaches.

You were talking to a colleague who mentioned that she has also done research on pane fruit. In her research, she found that pane fruit causes the release of neurotransmitter X and that neurotransmitter X causes headaches. You don't know much about neurotransmitter X because it is rarely found in the brain.

The second cover story dealt with the influence of vitamins on athletic performance.

Procedure

We manipulated the type of prevention within-subjects, so each participant read the two cover stories over the course of the experiment. The pairings between experimental condition and cover story were counterbalanced across participants.

Before beginning the experiment, participants were given some practice interpreting causal graphs. Participants were shown an example causal graph that involved simple causation, and the features of the graph were explained. Throughout the experiment, green arrows in causal graphs denoted generative causation and red arrows denoted preventative causation.

In the learning phase, participants first read the cover story. Then, before viewing any data, participants were shown a causal graph that summarized the information in the cover story and previewed a question that they would be asked later (see Figure 2 for an example). The graph showed the generative cause producing the mediating variable, the mediating variable producing the effect, and a node representing other causes producing the effect. A node for the preventer was shown with a question mark and no causal links, and participants were told that they would be asked to figure out how the preventer fit into the explanation.

Participants viewed data from four clinical trials for four different treatments: one where only the preventer was administered, one where only the generative cause was administered, one where both the generative cause and




Figure 2: Example graph provided to participants to summarize the information learned in the cover story and to preview a question about how the preventer (asmine juice) fit into the explanation

When the next group came in, you asked them how many had headaches and found:



Then you had them eat pane fruit, so this group:

- ate pane fruit.
- did NOT drink asmine juice.

After an hour, you found:



[next](#)

Figure 3: An example of a clinical trial. Each person in the clinical trial is represented by a cartoon face. The type of cartoon face indicates whether or not the person exhibits the effect.

preventer were administered, and one where neither the generative cause nor the preventer were administered. As shown in Figure 3, the data were presented through displays containing cartoon faces (following Buehner, Cheng, & Clifford, 2003). Each cartoon face represented a person in the clinical trial, and the type of cartoon face (happy face or sad face) indicated whether the person exhibited the effect. At the beginning of each clinical trial, 10 out of 50 people entering the trial exhibited the effect. Then the people in the trial were given a treatment, and the experiment showed how many people exhibited the effect at the end of the trial. The number of people exhibiting the effect at the end of the clinical trial depended on which candidate causes were present and on the experimental condition in accordance with Table 2.

Both before and after viewing the data, participants were told that the results had been replicated in much larger studies, and that they should consider any differences in the frequency of the effect to be reliable.

Participants were also provided with a summary of the clinical trials. The summary showed the frequency of the effect at the end of each clinical trial. The participants were encouraged to refer back to the summary as much as necessary.

In the inference phase, we measured the participants' beliefs about the candidate causes with counterfactual questions. One question assessed the causal power of the generative cause. For this question, participants were asked to suppose that there were 100 people who were not exposed to the generative cause or the preventer and who did not exhibit the effect. They were then asked how many of those people would have exhibited the effect had they been exposed to the generative cause.

To assess inferences about preventative scope, we constructed three preventative counterfactuals. Each counterfactual assessed the effectiveness of the preventer among a group of 100 people who exhibited the effect, but the likely cause of the effect varied between questions. In the *prevent|known* counterfactual, the group had been exposed to the known generative cause (e.g., "Suppose there are 100 people who ate pane fruit and have headaches"). In

the *prevent|unknown* counterfactual, the group had NOT been exposed to the generative cause (e.g., "Suppose there are 100 people who did NOT eat pane fruit but who have headaches"). Finally, in the *prevent|novel* counterfactual, the group had been exposed to a novel cause of the effect (e.g., "Suppose there are 100 people who recently stopped drinking coffee. They have been experiencing caffeine withdrawal and have headaches."). In each case, participants were asked to predict how many of the group would have exhibited the effect if they had also been exposed to the preventer (e.g., "If they had [ALSO] drank asmine juice, how many of them still would have had headaches?"). Although the *prevent|known* and *prevent|novel* questions left open the possibility that some of the effects were due to unknown or unmentioned causes, it is likely that many of the effects were due to the mentioned generative cause.

Finally, participants were asked to choose between two causal explanations in order to explain the observed data. The choice was presented as a choice between two causal graphs (see Figure 4). In one causal explanation, the preventer directly reduced the likelihood of the effect. In the other causal explanation, the preventer reduced the likelihood of the effect indirectly by preventing the mediating variable. Participants were also asked to explain why they chose the graph that they chose. These explanations were primarily intended to encourage reflection and were not formally analyzed.

Results

As expected, the answers to the counterfactual regarding the generative cause were similar across conditions. Participants in the broad prevention and narrow prevention conditions expected the generative cause to produce the effect for an average of 63.5 (SD=20.2) and 63.9 (SD=23.9) people respectively.

For the preventative counterfactual questions, participants were asked to estimate the number of cases where the effect still would have been present even if the group had been exposed to the preventer. By subtracting a participant's answer from 100 (i.e., the number of people in the group who had exhibited the effect), we obtained the participant's




Figure 4: The two causal graphs that participants were asked to choose between.

estimate of the number of cases that would have been prevented by the preventer. The results of these calculations are shown in Figure 5 for each of the preventative counterfactuals. When the effect occurred in the presence of the generative cause, participants in both conditions expected the preventer to reduce the occurrence of the effect. However, when the effect occurred in the absence of the generative cause or when the effect occurred due to a novel cause, the broad preventer was expected to be more influential than a narrow preventer.

The choices for the causal graphs are shown in Figure 6. Participants explaining narrow prevention were more likely to select the causal explanation where the preventer acted upon the mediating variable.

A multivariate ANOVA was performed with prevention condition (narrow or broad) as the independent variable. The dependent variables were the responses to the counterfactuals (known generative, prevent|known, prevent|unknown, and prevent|novel) and the forced choice between the graphical explanations. As expected, there was no effect




Figure 5: Participant estimates of the number of headaches that would have been prevented had the group been exposed to the preventer.

of prevention condition on the generative counterfactual, $F<1$, $p=.82$. The experimental condition also did not significantly influence responses to the prevent|known counterfactual, $F<1$, $p=.83$. Since the causal power of the preventer was higher in the narrow prevention condition than in the broad condition (.625 vs .5), one might have expected the narrow preventer to be more effective than the broad preventer when the generative cause is present. The predicted difference was relatively small, however, so there might have been too much noise in the data to detect it.

Statistical tests also confirmed that participants predicted more prevention in the broad prevention condition for prevent|unknown counterfactual, $F(1,78)=62.61$, $p<.001$ and prevent|novel counterfactual, $F(1,78)=63.69$, $p<.001$. Additionally, the increased preference for the graph with mediation in the narrow prevention condition was significant, $F(1,78)=6.49$, $p<.05$.




Figure 6: Preferred causal graph. This shows the proportion of participants who chose the explanation where the preventer stopped the mediating variable over the explanation where the preventer acted directly on the effect.

Discussion

The answers to counterfactual questions indicated that participants distinguished between broad and narrow prevention. Most notably, broad preventers were more likely to be generalized to situations where the generative cause was absent but where the effect occurred for some other reason. This pattern was found both when the cause of the effect was unknown and when the likely cause of the effect was a novel generative cause. Despite only observing one generative cause of the effect, participants distinguished between narrow and broad prevention by observing the influence of the preventer in the absence of the generative cause.

The graph choices suggest that participants believe that narrow prevention and broad prevention are associated with different causal explanations. When shown data consistent with narrow prevention, participants were more likely to endorse the causal explanation that involved mediation. However, the exact nature of this preference is unclear. Since participants were forced to choose between two explanations, their preference could be interpreted as either an endorsement of the chosen explanation or a rejection of the other explanation.

Furthermore, there are other explanations that were not considered in the current experiment. For instance, the narrow prevention data can be produced when the combination of the preventer and the generative cause is treated as a conjunctive preventer of the effect. There are also other explanations involving unobserved variables that were not considered here. Additional research is needed to describe people's preferences more completely.

Still, it is also worth noting that the causal explanation with mediation contained a causal relationship that was not directly supported by any evidence (i.e., the preventer preventing the mediating variable). Despite this, a strong majority of participants in the narrow prevention endorsed that explanation. This raises the possibility that people make inferences about unobserved mediating variables after observing narrow prevention. Other studies have identified related inferences about unobserved or hidden causes (e.g., Hagmayer & Waldmann, 2007; Luhmann & Ahn, 2007; Saxe, Tenenbaum, & Carey, 2005; Schulz & Sommerville,

2006). These inferences are usually interpreted as reflections of causal assumptions, such as the assumption that every effect has a cause. Although additional research is needed, inferences about a mediating variable may reflect causal assumptions about the nature of narrow prevention. Another possibility is that these assumptions reflect a preference for simple explanations (Lombrozo, 2007).

In conclusion, people can infer preventative scope from observations, and preventative scope influences further inferences. By default, narrow prevention is not generalized to circumstances where the effect is produced by novel causes. Finally, narrow and broad prevention are differentially compatible with different causal explanations.

References

Buehner, M. J., Cheng, P.W., & Clifford, D. (2003). Covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 29* (6). 1119-1140.

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

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction, *Cognitive Psychology, 51*, 334-384.

Hagmayer, Y., & Waldmann, M.R. (2007). Inferences about unobserved causes in human contingency learning. *Quarterly Journal of Experimental Psychology, 60* (3), 330-355.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology, 55* (3), 232-257.

Luhmann, C. & Ahn, W.-k. (2007). BUCKLE: A model of unobserved cause learning. *Psychological Review, 92* (3), 657-677.

Novick, L.R., & Cheng, P.W. (2004). Assessing interactive causal influence. *Psychological Review, 111*, 455-485.

Saxe, R., Tenenbaum, J.B., & Carey (2005). Secret agents: inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science, 16* (12), 995-1001.

Schulz, L.E., & Sommerville, J. (2006). God does not play dice: Causal determinism and children's inferences about unobserved causes. *Child Development, 77* (8), 427-442.