An information integration account for causal generalizations in uncertain causal networks

Moyun Wang

School of Psychology, Shaanxi Normal University, Xi'an 710062, China

Corresponding author: Moyun Wang, E-mail: wangmoyun@snnu.edu.cn

Abstract

In reasoning about common cause networks, given that a cause generates an effect, people often need to infer how likely the cause generate another effect. This causal generalization question has not systematically been investigated in previous research. We propose the information integration account for causal generalizations in uncertain casual networks with dichotomized continuous variables. It predicts that causal generalization is the joint function of conditional probabilities of causal links and cause strength indicated by the proportion of present collateral effects. Two experiments investigated causal generalizations in uncertain causal networks with and without probability distributions, respectively. It was found that in the presence of probability distributions there was the joint effect of conditional probability and cause strength on causal generalization; in the absence of probability distributions causal generalization depend only on cause strength. The overall response pattern favors the information integration account over the other alternative accounts.

Keywords: Causal Bayes networks; causal generalization; conditional probability; cause strength; the information integration account

Causal reasoning is pervasive in everyday life and scientific contexts. In causal reasoning, a current main question is how people infer a new causal link (e. g., the collision of two cars causes a quarrel between two car owners) from a known causal link (e. g., a collision of two cars causes a car damage). Such causal inference is causal generalization from a known causal link ($C \rightarrow E_1$) to another new causal link ($C \rightarrow E_2$). In such causal generalizations, the two causal links share a cause component, but differ in the effect component, and so both links form a common cause network.

Causal inferences/generalizations are based on causal knowledge. Causal probability knowledge represents causal dependencies and functional relations from causes to effects. The Causal Bayes network represents this kind of knowledge. Causal Bayes networks (henceforth referred to as “causal networks”) provide normative calculations for inferences on causal networks (Glymour, 2001; Gopnik, et al, 2004; Griffiths et al. 2008; Griffiths & Tenenbaum, 2009; Holyoak & Cheng, 2011; Lucas & Griffiths, 2010; Pearl, 2000; Rottman & Hastie, 2013). Causal networks represent the world as a probability distribution over all relevant variables by using a directed acyclic graph. A directed acyclic graph consists of a set of nodes and directed edges. A node represents a variable. A directed edge represents a direct causal relation from a cause to an effect, and the strength of the causal relation is represented by the conditional probability that the effect occurs given that the cause occurs. Most of the relevant psychological literature on causal networks focus on binary variables.

Causal networks have three main assumptions. The causal sufficiency assumption states that graphs must include all variables responsible for observed statistical information. The minimality assumption defines the simplest causal graphs as those with the fewest arrows while capturing all the dependencies in the data (Pearl, 2000). The Markov Assumption is the core assumption of

normative causal inference. It states that a given node, conditional on all its direct causes, is statistically independent of all other nodes except its direct or indirect effects (Glymour, 2001; Hausman & Woodward, 1999; Pearl, 2000; Spirtes et al., 2001; Rottman & Hastie, 2013). This means that when its direct causes are certain, the given node is statistically independent of all other nodes except its direct or indirect effects. The certain states of the direct causes of the given node screen off their direct effects from other variables. This is the screening-off rule. The direct causes are screening-off variables. When its direct causes are not constant, the given node is no longer statistically independent of all other nodes except its direct or indirect effects. The screening-off rule is based on the basic assumption that an effect variable is the function of its direct cause variables and so when the direct cause variables remain constant, the probability of the effect variable accordingly remains constant regardless of all other variables.

The screening-off rule is typically illustrated by the common cause and chain networks in Figure 1 where each node denotes a binary variable with two alternative states (positive versus negative states). In the common cause network, when the state of node C as the common cause is constant, effect nodes E_1 and E_2 are independent from each other. Thus, $P(E_2|C, E_1)$, the probability that E_2 occurs given that C and E_1 are present, should be equal to $P(E_2|C, \neg E_1)$, the probability that E_2 occurs given that C is present but E_1 is absent. Here, the two probability inferences concern the causal generalization question whether causal link $C \rightarrow E_2$ can follow from causal link $C \rightarrow E_1$. In the chain network, when the state of middle node M is certain, node E is independent from node C. Current research focuses on the common cause network because it involves typical causal generalizations between causal links sharing a common cause.

Figure 1. Two typical causal networks used to test the screening-off rule

Most psychological studies testing the screening-off rule used binary variables as node variables in causal networks. These studies found that human causal inferences violate the screening-off rule in causal networks with binary variables (Burnett, 2004; Mayrhofer et al., 2010; Park & Sloman, 2013; Rehder & Burnett, 2005; Rehder, 2014, 2018; Rehder & Waldmann, 2017; Rottman & Hastie, 2013, 2016; Walsh & Sloman, 2004, 2008). In the common cause network, when the state of cause node C is certain, people tend to infer the unknown effect node from one known effect node in a way of positive correlation. In the chain network, when the state of middle node M is certain, people still use known cause node C to infer unknown effect node E in a way of positive correlation. When the direct cause of a given effect is absent, some studies found the violation of the screening-off rule, whereas some other studies did not find it. When the direct cause of a given effect node is present, most studies found the violation (Mayrhofer, et al., 2008, 2010; Park & Sloman, 2013; Rehder, 2014). These findings show that in a common cause network, causal inferences for an unknown effect node are sensitive to the states of its collateral effect nodes.

However, in many of these studies (Park & Sloman, 2013; Rottman & Hastie, 2016; Rehder, 2014; Rehder & Waldmann, 2017), violations of the screening-off rule in causal networks with binary variables may be false, because so called binary variables used by them are essentially

dichotomized continuous variables where at least one of two categories in a binary variable can vary in magnitude, but not real binary variables such as gender (males versus females) where each category does not vary in magnitude. For example, Rehder (2014) used a common cause network where high interest rates cause both large trade deficits and low retirement savings, and a forced-choice procedure where participants were presented with a pair of cases, and asked to choose which was more likely to possess a particular variable value. In the causal network, the binary variables were the dichotomies of the continuous variables: interest rates (low or high), trade deficits (small or large), and retirement savings (low or high). Each category in each variable can vary. For example, high interest rates can vary from 8% to 15%. These continuous variables were not real binary category variables. For a pair of causal cases with the above common cause network, the list of variables for one economy was “high interest rates,” “large trade deficits,” and “???” (indicating that the value for the third variable, retirement savings, was unknown), those for the second economy might be “high interest rates,” “???” and “???,” and participants were asked which economy was more likely to have low retirement savings. Some subjects chose the first economy. This violated the screening-off rule, according to which the two economies were equally likely in that the common cause “high interest rates” was present in the two economies. However, this violation might be false. These subjects might consider that “high interest rates” as the cause of “large trade deficits” in the first economy was higher than “high interest rates” as the cause of an unknown effect in the second economy, because a known effect “large trade deficits” suggested higher “high interest rates”, compared with an unknown effect. Thus, they judged that higher “high interest rates” in the first economy was more likely to cause the collateral effect “low retirement savings,” compared with the second economy. These subjects may have perceived that the common cause differed in magnitude between the two economies. When perceived cause strength varied,

these subjects' inferences did not essentially violate the screening-off rule. In sum, violations of the screening-off rule in Rehder's experiments (2014) might be a false phenomenon resulted from magnitude variation of dichotomized continuous variables. This is suggested by the "uncertainty model" that uncertainty concerning the exact states of binary variables as dichotomies of continuous variables could produce apparent Markov Violations (Rehder & Burnett, 2005; Rottman & Hastie, 2016).

A causal network with dichotomized continuous variables is essentially an uncertain network in that the simple dichotomy of a continuous variable is unable to determine the exact value of each category in the dichotomy. Many previous studies testing the screening-off rule examined causal inferences in such uncertain networks. In the above example, causal reasoning about the uncertain causal network is essentially causal generalization from a causal link (high interest rates lead to large trade deficits) to another causal link (high interest rates lead to low retirement savings). Current research aims to investigate how people make causal generalizations in uncertain causal networks. For causal generalization about uncertain causal networks, there is an intuitive correlation principle: when all other conditions are equal, a stronger cause can generally generate more effects than a weaker cause; conversely, a cause that generate more effects than another cause is generally stronger than the latter cause. This principle represents the positive correlation between cause strength and the number of effects. We can call it the strength-number correlation principle. For two comparable causes that are at different levels of a continuous variable as the common cause in a common cause network, one cause is stronger and the other is weaker. For example, for two-cars collisions, there is a binary variable: the presence or absence of collisions. When collisions occur, collision force is a continuous variable and so may differ in magnitude and so differ in the number

of effects. In particular, a stronger collision can generally generate more effects (such as car damage, car firing, a traffic jam, etc.) than a weaker collision (with only the effects of car damage and a traffic jam). From the strength-number correlation principle, it follows the causal generalization principle that if a stronger cause can generate more known effects than a weaker cause, it will be more likely to generate an unknown effect than the latter; and if a cause generates more known effects than another cause, it will be stronger than the latter, and so it will be more likely to generate an unknown effect than the latter cause. In the above example, the stronger collision with more effects is more likely to generate a quarrel between two car owners than the weaker collision with fewer effects. Conversely, the collision with more effects (such as car damage, car firing, and a traffic jam) is stronger than the collision with fewer effects (only the effects of car damage and a traffic jam), and so is more likely to generate a quarrel between two car owners than the latter. In the causal generalization principle, the number of known effects of a cause indicates the strength of the cause, which further indicates the likelihood of an unknown effect to be inferred.

The causal generalization principle treats the information of collateral effect nodes as the indicator of cause strength, and so can explain the previous finding of violations of the screening-off rule in uncertain causal networks with false binary variable. Moreover, prior studies have found that in common cause networks, human causal inferences are sensitive to changes of probability distributions (including conditional probabilities of causal links) over common cause networks (Rottman & Hastie, 2013, 2016). Overall, in common cause networks, causal inferences are sensitive to both the information of collateral effect nodes and conditional probabilities of causal links. Thus, causal generalizations will be the joint function of cause strength and conditional

probability. Accordingly, we integrate the causal generalization principle with conditional probability, and propose an information integration account for causal generalizations that incorporates the two kinds of information. In the information integration account, causal generalization is the joint function of cause strength and conditional probability when both kinds of information are available, and so there will be the joint effect of cause strength and conditional probability on causal generalization. When one of the two factor is constant or absent, causal generalization will depend on the other.

Current research aims to investigate how people generalize from some known causal links to an unknown causal link in uncertain causal networks, but not to test the screening-off rule in causal networks with real binary variables. In real world and everyday life, people often represent causal networks as uncertain or vague, but not normative. Normative causal networks may be the ideal construction of experts or researchers. Thus, uncertain causal networks are ecologically more valid, and so are more worth investigating.

Next, we specify main alternative accounts for causal generalization in uncertain causal networks. Because causal inferences that some previous studies on the screening-off rule investigated were essentially causal generalizations about uncertain causal networks with dichotomized continuous variables, some previous accounts for violations of the screening-off rule may be the alternative accounts for causal generalization in uncertain networks. There are six possible accounts for violations of the screening-off rule in uncertain networks (Rottman & Hastie, 2016). They fall into four categories: the association account, the heuristic account, the underlying mechanism account, and the Beta inference account. We introduce these accounts as follow.

The first category is the association account that the violation of the screening-off rule is caused

by associative reasoning based on processing similarity between causal cases and prototypes (Darlow & Sloman, 2010; Evans, 2008; Park & Sloman, 2013; Rehder, 2014). Rehder (2014) claimed that some subjects often reasoned “associatively,” that is, assumed that the presence of one positive node implied the presence of other positive nodes that are causally related, even those that were (according to the Markov condition) conditionally independent. The essence of such association is that if a causal case has more known positive nodes similar to positive nodes in the causal network prototype, it is more likely to have an unknown positive nodes similar to the corresponding positive node in the causal network prototype. However, there is no direct evidence that the violation of the screening-off rule is caused by associative reasoning. Moreover, the association account can not explain Park and Sloman’s finding (2013) that when the strength of common and middle causes remained at zero, that is, in the negative state of screening-off binary variables, causal inferences for an unknown effect remained stable regardless of the states of the other causally related variables, not violating the screening-off rule.

The second category is the heuristic account (Rottman & Hastie, 2016). The idea is that the presence (absence) of other causally related positive nodes is the diagnostic cue of the presence (absence) of the inferred positive node regardless of the state of screening-off nodes, and so more present (absent) positive nodes imply higher (lower) probability judgments for the inferred positive node, and thereby the violation occurs (Hastie & Dawes, 2010; Hogarth & Karelaia, 2007; Karelaia & Hogarth, 2008; Park & Sloman, 2013; Rottman & Hastie, 2016). Park & Sloman (2013) found that causal inferences were not consistent with the prediction of the heuristic account. However, Rottman & Hastie (2016) demonstrates that the heuristic account can explain Markov violations on chain and common cause networks.

The underlying mechanism account assumes that causal scenarios presented to participants are insufficient and so participants imagine an additional unobserved variable or mechanism affecting the observed variables by appealing to prior causal knowledge (Park & Sloman, 2013; Rehder, 2014; Rottman & Hastie, 2013, 2016; Walsh & Sloman, 2008). However, there is little direct evidence that people can report such unobserved underlying mechanisms (Rehder, 2014; Rottman & Hastie, 2013). In many situations, internal causal mechanisms are often unobservable and unavailable. Thus, causal inferences cannot often depend on such unavailable underlying mechanisms.

Park & Sloman (2013) manipulated whether different causal links in a common cause structure are supported by the same or different mechanisms. They found that in the positive state of screening-off binary variables, the violation of the screening-off rule was larger when different causal links shared the same mechanism than when they had different mechanisms. This suggests that in common cause structures, people are more likely to make causal generalizations between causal links with the shared mechanism than between causal links with different mechanism. We call this the shared mechanism account for causal generalizations.

The Beta inference account claims that when a causal network is supplemented with trial-by-trial learning data from experience, it is possible to make judgments directly from the data itself (Rottman & Hastie, 2016). For example, when making the judgment $P(x_1 = 1 | x_2 = 1)$ in a common cause network including a common cause (y) and two direct effects (x_1 and x_2), a learner could simply divide the number of trials in which $x_1 = 1$ and $x_2 = 1$ by the number of trials in which $x_2 = 1$. These judgments would be exactly the same as the more typical version of the CBN theory in which learners infer parameters from learning data, and then uses parameters to make judgments.

When a causal network is supplemented with learning data in the summary format, there is no direct data to calculate $P(x_1 = 1|x_2 = 1)$. Thus, the Beta inference account does not apply to causal networks with learning data in the summary format.

In summary, in uncertain causal networks, only the information integration account takes into account the joint influence of cause strength and conditional probability on causal generalization. The other alternative accounts all neglect the influence of cause strength. The present experiments aimed to investigate whether cause strength affects causal generalization in uncertain common cause networks and to distinguish between the information integration account and the other alternative accounts.

Overview of the present experiments

The present experiments investigated whether causal generalization would vary with perceived cause strength in uncertain common cause networks where each node had two alternative states (the positive versus negative state of the event) in which the positive state can vary in magnitude. Experiment 1 investigated uncertain causal network prototypes with probability distributions. Experiment 2 investigated uncertain causal network prototypes without probability distributions. In the absence of probability distributions, a causal network is more uncertain in that conditional probabilities of causal relations in the network are unknown and uncertain.

The present experiments investigated only causal networks with probabilistic causal relations, but not causal networks with deterministic causal relations, because deterministic causal relations imply that a cause necessarily bring about its effects and so it is impossible that some effects of the cause does not occur. Moreover, because Park and Sloman (2013) have found that in negative states of screening-off variables, causal inferences did not violate the screening-off rule, Experiment 1 limits

itself to positive states of screening-off variables.

Experiment 1

Method

Participants

70 (30 males, 40 females) undergraduate students from Shaanxi Normal University in China participated in the experiment. The two experiments received ethical approval from the University Human Research Ethics Committee, and all of the participants signed a written consent form.

Design and Materials

The experiment used a paper-pen questionnaire test in which there were two between-subjects problems with different causal scenarios (see the experimental materials). The two problems had the same design shown in Table. In each problem, participants learned causal network prototypes with the summary presentation format, and then made causal inferences for some causal cases. The summary presentation format could reduce working memory load and so participants can clearly observe all information of causal network prototypes in causal problems. Moreover, the summary presentation format rendered the Beta inference account for trial-by-trial causal learning not applicable to the present experiment.

Each problem involved two parts: a causal network prototype and some causal cases. Causal network prototypes involved some causal links and provided the conditional probability of each causal link, that is, the conditional probability that an effect occurs given that the cause occurs. In a causal network prototype, there were typical and atypical causal links. A typical causal links had the conditional probability of 75% indicating the strong causal link that generally a cause can generate the effect. Two kinds of atypical causal links had the conditional probability of 50% or 25%

indicating the moderate or weak causal link that generally a causes can not generate the effect.

These conditional probability settings conformed to the parameter independence assumption in causal Bayes networks that conditional probabilities determined by the Markov factorization do not depend on one another. In the part of causal cases, for each causal case, participants were asked to infer the strength of the screening-off variable and the probability of the positive state of the unknown node. Unknown nodes to be inferred were target nodes. Other known nodes were non-target nodes.

For common cause networks, most of the previous studies investigated common causes with two direct effects. This common cause network with minimal direct effects is particular. In real situations, a common cause often generates more than two direct effects. For example, an earthquake can cause many direct effects. The present experiment used common cause networks with five direct effects that are ecologically more valid. Table 1 shows the common cause network where node C is the common cause and nodes A, B, D, E, and F are five direct effects. The six nodes form five causal links in the network. The example problem is illustrated as Problem 1 in the materials.

This experiment would distinguish between five accounts for causal generalizations in uncertain causal networks: the information integration account, the similarity-based association account, the heuristic account, the shared mechanism account, and the conditional probability account in the normative causal Bayes network theory. The experimental design and theoretical predictors for probability judgments are presented in Table 1. In the common cause network problem in Table 1, there are three causal cases (see the example problem in the following materials). For the five accounts, there were four predictors (conditional probability, overall similarity, typical similarity,

cause strength) that were simultaneously manipulated. We specify these predictors and the five accounts as follow.

Table 1.

The problem design and settings of predictors for probability judgments in common cause networks

Prototype	Nodes						Overall	Typical	Cause	Conditional	
	C	A	B	D	E	F	similarity	similarity	strength	probability	
	+	75%	25%	75%	75%	50%	Nodes E/F	Node E	Nodes E/F	Node E	Node F
Case 1	+	-	+	-	E ₁ ?	F ₁ ?	0/3	0/2	1/3	75%	50%
Case 2	+	+	+	+	E ₂ ?	F ₂ ?	2/3	2/2	3/3	75%	50%
Case 3	+	+	-	-	E ₃ ?	F ₃ ?	2/3	1/2	1/3	75%	50%

Notes: + denotes the positive state (the presence of causes or effects) of nodes. - denotes the negative state (the absence of causes or effects) of nodes respectively. ? denotes unknown target nodes to be inferred. The same below. In the prototype, the probability below each of nodes A, B, D, E, and F denotes the conditional probability that the corresponding effect occurs given that cause C occurs.

The conditional probability account predicted that for the three causal cases probability judgments would ordinally accord with conditional probability indicated by the corresponding conditional probability in the causal network prototype. Thus, in each causal case, probability judgements would be higher for node E than for node F. For each of nodes E and F, probability judgements would not differ between the three causal cases.

The information integration account predicted that when conditional probability remained constant, cause strength judgments and probability judgments both would ordinally accord with cause strength indicated by the proportion of the present effects in a causal case; when cause strength remained constant, probability judgments would ordinally accord with conditional

probability. Thus, for each response variable of each target nodes, there would be case 2 > case 1 = case 3. For each causal case, probability judgments would be higher for node E than for node F.

The heuristic account predicted that probability judgments for each target node would increase with the proportion of present effects in causal cases, and so there would also be case 2 > case 1 = case 3 for probability judgments for each target node. It predicted that in each causal case, there would be no differences between target nodes E and F because there was the same proportion of present effects for both target nodes. It did not predict cause strength judgments.

The similarity-based association account predicted that probability judgments would ordinally accord with overall similarity indicated by the proportion of the matching causal links between a causal case and the causal network prototype. A causal case with more causal links matching the corresponding causal links in the causal network prototype was more similar to the causal network prototype. For example, the common cause prototype in problem 1 in the materials had five causal links with different conditional probabilities (25%, 50%, and 75%). In Case 3 one strong causal link ($C \rightarrow A$) occurred (+) and one weak causal links ($C \rightarrow B$) occurred (-). In Case 3, among the three known causal links, there were two causal links matched the corresponding causal links in the causal network prototype. Thus, the overall similarity between Case 3 and the causal network prototype was 2/3. Here, for causal networks, it is appropriate that similarity was measured in terms of structure (causal links) similarity but not component (nodes) similarity, because it is structure similarity, rather than component similarity, that can reflect causal structures/relations in causal networks.

For the shared mechanism account, in the common cause network prototype in Table 1, three typical causal links with conditional probability 75% shared a strong causal connection mechanism, but two atypical causal link had a moderate or weak causal connection mechanism. Thus, the three

conditional probabilities corresponded to three different causal connection mechanisms. The causal link $C \rightarrow E$ (corresponding to the unknown node) shared the same strong causal connection mechanism as the other two typical causal links ($C \rightarrow A$, $C \rightarrow D$). Thus, for the shared mechanism account, the predictor was typical similarity indicated by the proportion of the matching typical causal links of a causal case and a network prototype. For example, the causal network prototype of problem 1 had three typical causal links sharing a strong causal connection mechanism. In Case 3, among the two known typical causal links ($C \rightarrow A$, $C \rightarrow D$), only one matches the typical causal link in the prototype. Thus, the typical similarity between Case 3 and the causal network prototype was $1/2$. The shared mechanism account predicted that probability judgments for node E would ordinarily accord with typical similarity. For node F, causal link $C \rightarrow F$ had the moderate causal connection mechanism different from the causal connection mechanism in the other causal links. Thus, probability judgments for node F would not vary with the other causal links and so would not differ between the three causal cases.

In summary, the five accounts made different predictions for the size order of probability judgments. Only the information integration account predicted that cause strength judgments would be accord with the proportion of present effects in causal cases, and cause strength and conditional probability would both affect probability judgments. Thus, the present design was able to distinguish between the five accounts.

Experimental materials were written questionnaires containing a problem on a sheet. We designed two similar problems with different causal scenarios. One causal scenario was a neuron network, and the other was that a bomb exploding causes bottles to tip over. Participants lacked background knowledge about neuron networks. But they probably had knowledge about bomb

exploding and domino networks. The two kinds of different materials were representative for investigating causal generalization in uncertain networks.

The original Chinese questionnaire of one example problem is translated into the following English version. The other problem is presented in the appendix.

Problem

Six neurons constitute a network in which five independent neurons (A, B, D, E and F) have a common cause neuron C. Neurons A, B, D, E and F have independent and direct neural connections with C respectively, as shown by the arrows in the figure below. There is no connection among neurons A, B, D, E and F. Neurons have the two states of firing (+) and not firing (-). For each 100 times that neuron C fires, the probabilities that neuron C excites the other neurons are represented under the corresponding neurons respectively. If neuron C doesn't fire, then neurons A, B, D, E and F will not fire.

Please answer the following questions and fill your ratings in brackets. Strength ratings should be given on a scale from 1-7 where larger values indicate greater strength. Probability ratings should be given on a scale from 0% (impossible) to 100% (definite). The greater the rating is, the greater the probability of the occurrence of firing is.

(1) Given a firing of neuron C and its effects as shown in the figure below, what is the firing strength of neuron C?

() What is the probability that neuron E will fire? () What is the probability that neuron F will fire? ()

- (2) Given a firing of neuron C and its effects as shown in the figure below, what is the firing strength of neuron C?
- () What is the probability that neuron E will fire? () What is the probability that neuron F will fire? ()

- (3) Given a firing of neuron C and its effects as shown in the figure below, what is the firing strength of neuron C?
- () What is the probability that neuron E will fire? () What is the probability that neuron F will fire? ()

The two problems were balanced across participants. The presentation order of causal cases within each problem was counterbalanced. One was the order in the example problem, and the other was the reversion of the former.

Procedure

Participants voluntarily participated in the experiment. The experiment was conducted in a large classroom. They took about eight minutes to complete the questionnaires.

Results

The two problems with different causal scenarios did not differ in performance, and so we collapsed the results of the two problems. The overall results are presented in Table 2.

Table 2

The results M (SD) in Experiment 1

		Case 1	Case 3	Case 2
Cause strength judgments		2.93 (1.57)	3.09 (1.22)	5.43 (1.33)
Probability	Node E	.45 (.28)	.46 (.24)	.66 (.24)
judgments	Node F	.37 (.20)	.39 (.20)	.51 (.22)

For cause strength judgments, we conducted a paired sample t -test for each pair of the three causal cases. Cause strength judgments were higher in case 2 than in cases 1 and 3 [$t(69) = 10.79, p < .001$; $t(69) = 11.35, p < .001$]. Cause strength judgments showed no differences between case 2 and 3. Thus, cause strength judgments ordinally accorded with the proportion of present effects in causal cases. This confirms the strength-number correlation principle that the proportion of present effects indicates perceived cause strength. Thus, the strength manipulation was valid.

For probability judgments, because the four predictors were simultaneously manipulated, for each pair of the three causal cases we performed a 2 (the paired two cases) $\times$ 2 (node E with conditional probability 75% vs. node F with conditional probability 50%) repeated ANOVA. For the pair of cases 2 and 1, it revealed that probability judgments were higher in case 2 than in case 1 [$F(1, 69) = 23.10, MSE = 2.08, p < .001, \eta^2 = .25$], and probability judgments were higher for node E than for node F [$F(1, 69) = 33.84, MSE = 0.84, p < .001, \eta^2 = .33$], and there was a significant interaction between the two factors [$F(1, 69) = 5.67, MSE = 0.09, p < .05, \eta^2 = .08$]. For the pair of cases 2 and 3, it revealed that probability judgments were higher in case 2 than in case 3 [$F(1, 69) = 20.84, MSE = 1.74, p < .001, \eta^2 = .23$], and probability judgments were higher for node E than for node F [$F(1, 69) = 31.72, MSE = 0.80, p < .001, \eta^2 = .32$], and there was a significant interaction

between the two factors [$F(1, 69) = 6.37$, $MSE = 0.11$, $p < .05$, $\eta^2 = .08$]. For the pair of cases 1 and 3, it revealed that probability judgments showed no differences between cases 1 and 3, and probability judgments were higher for node E than for node F [$F(1, 69) = 15.67$, $MSE = 0.34$, $p < .001$, $\eta^2 = .18$]. There was no significant interaction between the two factors. Overall, for probability judgments for each target node, there was Case 2 > Case1 $\approx$ Case 3, showing the cause strength effect. This result is consistent with only the prediction of the information integration account and the heuristic accounts. In each causal case, probability judgments were higher for node E than for node F, showing the conditional probability effect. This result is consistent with the prediction of the information integration account and the conditional probability account. Thus, probability judgments for target nodes was the joint effect of conditional probability and cause strength. This result confirms the prediction of the information integration account that probability judgments would be the joint function of conditional probability and cause strength. The overall pattern of probability judgments favors the information integration account over the other accounts.

Experiment 2

Method

Participants

48 (20 males, 28 females) undergraduate students from Shaanxi Normal University in china participated in the experiment.

Design and Materials

Experiment 2 used a paper-pen questionnaire test to investigate causal generalizations in common cause network prototypes without probability distributions. In the questionnaire, there were two between-subjects problems with different causal scenarios. The two problems had the same causal

scenarios as those in Experiment 1, except that the design differed from that in Experiment 1. We illustrate the two problems with the following neuron network problem translated from the original Chinese questionnaire. The other problem was the bomb problem. The two problem had the same design shown in Table 3.

Problem

Five neurons constitute a network in which four independent neurons (A, B, D, and E) have a common cause neuron C. Neurons A, B, D, and E have independent and direct neural connections with C respectively, as shown by the arrows in the figure below. There is no connection among neurons A, B, D, and E. Neurons have the two states of firing (+) and not firing (-). For each 100 times that neuron C fires, the probabilities that neuron C excites the other neurons are represented under the corresponding neurons respectively.

Please answer the following questions and fill your ratings in brackets. Strength ratings should be given on a scale from 1-7 where larger values indicate greater strength. Probability ratings should be given on a scale from 0% (impossible) to 100% (definite). The greater the rating is, the greater the probability of the occurrence of firing is.

- (1) Given a firing of neuron C and its effects as shown in the figure below, what is the firing strength of neuron C?
 () What is the probability that neuron E will fire? ()

- (2) Given neuron C not firing and its effects as shown in the figure below, what is the probability that neuron E

will fire? ()

(3) Given a firing of neuron C and its effects as shown in the figure below, what is the firing strength of neuron C?

() What is the probability that neuron E will fire? ()

The two problems were balanced across participants. In each problem, the presentation order of the three cases was counterbalanced. The first order was presented in the above questionnaire, and the other was the reversed version of the former.

Table 3.

The problem design and settings of predictors for probability judgments in Experiment 2

Causal Prototypes	Nodes					Cause strength
	C	A	B	D	E	
Case 1	+	+	+	-	E ₁ ?	2/3
Case 2	-	-	-	-	E ₂ ?	0/3
Case 3	+	-	+	-	E ₃ ?	1/3

In Table 3, cases 1 and 3 where the common cause occurred were target items, whereas case 2 where the common cause did not occur was one non-target item. In the absence of probability

distributions in causal network prototypes, there was only the predictor causa strength indicated by the proportion of present effects, but there were not the predictors: conditional probability, typical and overall similarity. Thus, the conditional probability account based on conditional probability, the shared mechanism account based on typical similarity, and the similarity-based association account based on overall similarity were not applicable to this experiment. The information integration account predicted that when conditional probability was absent, probability judgments would increase with cause strength. Thus, for both causa strength and probability judgments there would be Case 1 > Case 3 > Case 2. The heuristic account predicted the same size order of probability judgments as the information integration account. However, it did not predict cause strength judgments.

The two problems were balanced across participants. The presentation order of causal cases within each problem was counterbalanced. One was the order in the example problem, and the other was the reversion of the former.

Procedure

The procedure was similar to that in Experiment 1.

Results

The two problems with different causal scenarios did not differ in performance, and so we collapsed the results of the two problems. The overall results are presented in Table 4.

Table 4 The results *M* (*SD*) in Experiment 2

	Case 1	Case 2	Case 3
Cause strength judgments	4.46 (1.27)	-	3.29 (1.73)
Probability judgments	.59 (.20)	.08 (.15)	.42 (.27)

For cause strength judgments, we conducted a paired sample t -test for the difference between cases 1 and 3. The average strength judgment was higher in case 1 than that in case 3, $t(47) = 3.33$, $p < .01$. For probability judgments, we conducted a paired sample t -test for the difference between cases 1 and 3. The average probability judgment was higher in case 1 than that in case 3, $t(47) = 3.07$, $p < .01$. Cause strength judgments increased with the proportion of present effects. Probability judgments increased with cause strength judgments. The overall response pattern favors the information integration account over the heuristic account.

General discussion

Causal generalizations did not differ between the two problems (neuron vs. bomb materials). Participants generally lacked background knowledge of neuron networks. But they probably had background knowledge of bomb networks. Thus, the variable of background knowledge does not affect causal generalizations, and the response pattern were reliable. This finding suggests that the underlying mechanism account is not applicable to the present experiments.

In uncertain causal networks with probability distributions, probability judgments for target nodes was the joint effect of conditional probability and cause strength. In uncertain causal networks without probability distributions, probability judgments for target nodes depended on only cause strength. In both the presence and the absence of probability distributions, probability judgments for a target node increased with perceived cause strength, which increased with the proportion of present effects. This result does not violate the screening-off rule. In causal cases, probability judgments were essentially causal generalizations from known causal links to the unknown causal link. Thus, causal generalizations increased with perceived cause strength indicated

by the proportion of present effects. This result confirms the strength-number correlation principle and the causal generalization principle in the introduction. The overall response pattern of probability and cause strength judgments favors the information integration account over the heuristic account, the shared mechanism account, the similarity-based association account, and the conditional probability account.

In Experiment 1, the joint effect of conditional probability and cause strength on probability judgments confirms the prediction of the information integration account that probability judgments would be the joint function of conditional probability and cause strength. The finding of the joint effect favors only the information integration account. All the other alternative accounts fail to predict the joint effect. The joint effect distinguishes between the information integration account and the other alternative accounts.

The joint effect implies that causal inference for target nodes is a mixed process that involves both inductive and deductive reasoning. In particular, the joint effect involves the cause strength effect that probability judgments increase with perceived cause strength indicated by the proportion of present collateral effects in causa cases, and the conditional probability effect that probability judgments increase with conditional probability. The cause strength effect reflects inductive reasoning from present collateral effects to perceived cause strength and then probability judgments. The conditional probability effect reflects deductive reasoning from conditional probabilities in causal network prototypes to probability judgments for target nodes in causal cases. However, pure inductive or deductive reasoning is insufficient to explain the joint effect. Probability judgment for target nodes is the compromise result of inductive and deductive reasoning. Thus, in uncertain causal networks, causal generalization for target nodes is a mixed

process of inductive and deductive reasoning, but not pure inductive or deductive reasoning.

In an uncertain causal network, perceived cause strength increased with the proportion of present effects in causal cases, as confirming the strength-number correlation principle. This principle may be a heuristic of cause strength judgment. However, it differs from the heuristic account of probability judgments. The heuristic account predicts only probability judgments, without considering cause strength. Moreover, the heuristic account does not predict the joint effect. Thus, the heuristic account is unable to explain the overall response patterns of probability and cause strength judgments. Overall, the information integration account is essentially different from the heuristic account.

In Experiment 1, the results don't support the shared mechanism account. In the shared mechanism account, people are more likely to make causal generalizations between causal links sharing the same mechanism than between causal links with different mechanism. As revealed by the results of the present common cause network with three strong causal links sharing the same mechanism in Experiment 1, when among these three causal links some causal links occurred in causal cases but some other not, probability judgements didn't depend on the proportion of present strong causal links, but on the proportion of present effects regardless of whether the four known causal links shared the same mechanism. This result clearly favors the information integration account over the shared mechanism account. In the information integration account, causal links with positive effect nodes in causal cases indefinitely share the same mechanism, that is, they may share different mechanisms.

In Experiment 1 using uncertain common cause networks with dichotomies of continuous variables, when perceived cause strength remained stable, probability judgments remained stable.

This result conforms to the screening-off rule. When perceived cause strength varied, probability judgments varied with former. This result is not incompatible with the screening-off rule. The overall response pattern of probability judgments confirms our previous conjecture that in many of the previous studies violations of the screening-off rule may be false because they used false binary variables that were essentially continuous variables.

In Experiment 1, when the conditional probability of target node E remains 75%, probability judgments varied with the proportion of present effects in causal cases, violating the prediction of conditional probability that probability judgments should remain unchanged. Causal generalizations were sensitive to the relationship or association between known causal links and unknown causal links to be inferred. This result violates the parameter independence assumption that in normative causal Bayes networks, conditional probabilities determined by the Markov factorization do not depend on one another, and so probability judgments for target nodes in causal cases should not depend on states of non-target nodes. This violation may be due to that in an uncertain common cause network, common cause strength in a positive state is variable and uncertain, resulting in the correlation between target and non-target nodes. If common cause strength remains constant, probability judgments for target nodes will more accord with conditional probability. This conjecture is worth future investigation.

The present finding of violations of the parameter independence assumption suggests that the parameter independence assumption in causal network prototypes may not be applicable to individual causal cases. In individual causal cases, causal generalizations are sensitive to the relationship between known causal links and unknown causal links to be inferred and so no longer conform to the parameter independence assumption. The parameter independence assumption may

hold only at the level of causal network prototypes that represent general causal statistical laws. At the level of individual causal cases, general causal statistical laws may no longer be applicable. Future study could investigate whether the parameter independence assumption in causal network prototypes is applicable to individual causal cases.

In Experiment 2 using uncertain causal networks without probability distributions, the absence of probability distributions resulted in that probability judgments for target nodes depended on only cause strength indicated by the proportion of present effects. In this situation, causal generalization was pure inductive reasoning from known causal links to cause strength and then inferring unknown causal links, without considering the parameter independence assumption. This does not violate the parameter independence assumption postulated by normative causal Bayes networks, which is not applicable to uncertain causal networks without probability distributions.

In the information integration account, the strength-number correlation principle is congruent with the force theory of causal relations (Wolff, 2007; Wolff & Barbey, 2015). The force theory holds that people specify causal relations in terms of configurations of forces that are evaluated with respect to an endstate. In particular, it assumes that people specify the origin and direction of various forces in a configuration of forces, and to a limited extent, the magnitudes of these forces. Conceptually, a cause strength is similar to the magnitude of a force. Each cause has a force. Intuitively, a larger force will generate more direct effects, and more effects require a larger force as the cause. For example, in the bomb explosion problems, a stronger explosion can tip over more bottles, and more bottles tipping over requires a stronger explosion. Thus, the strength or force of a cause is directly related with the number of its direct effects. The strength-number correlation principle is an extension of reasoning about physical forces.

Conclusion

The present two experiments investigated causal generalizations in uncertain causal networks, and demonstrated the following findings. In the presence of probability distributions, causal generalizations are the joint function of conditional probability and cause strength. In the absence of probability distributions, causal generalizations depend only on cause strength. In both the presence and the absence of probability distributions, causal generalizations increase with perceived cause strengths, which increase with the proportion of present effects. The overall response pattern of probability and cause strength judgments favors the information integration account over the other alternative accounts.

References

- Burnett, R. C. (2004). Inference from complex causal models (Doctoral dissertation). Retrieved from ProQuest Dissertations and Theses. (UMI No. 3156566)
- Darlow, A., & Sloman, S. A. (2010). Two systems of reasoning: Architecture and relation to emotion. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 382–392.
- Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–278.
- Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- 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, 3–32.
- Griffiths, T. L., Kemp, C., Tenenbaum, J. B. (2008). *Bayesian models of cognition*. In *Cambridge*

- Handbook Computational Psychology*, ed. R. Sun, pp. 59–100. London: Cambridge Univ. Press
- 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, 661–716.
- Hastie, R., & Dawes, R. M. (2010). *Rational choice in an uncertain world*. Thousand Oaks, CA: Sage.
- Hausman, D. M., & Woodward, J. (1999). Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science*, 50, 521–583.
- Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, 114, 733–758.
- Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inference as a rational process: The new synthesis. *Annual Review of Psychology*, 62, 135–163.
- Karelaia, N., & Hogarth, R. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin*, 134, 404–426.
- Mayrhofer, R., Goodman, N. D., Waldmann, M. R., & Tenenbaum, J. B. (2008). Structured correlation from the causal background. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the thirtieth annual conference of the cognitive science society* (pp. 303–308). Austin, TX: Cognitive Science Society.
- Osman, M. (2004). An evaluation of dual-process theories of reasoning. *Psychonomic Bulletin & Review*, 11, 988–1010.
- Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the Markov property in causal reasoning. *Cognitive Psychology*, 67, 186–216.

- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, England: Cambridge University Press.
- Rehder, B. (2011). Reasoning with conjunctive causes. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 1406–1411). Austin, TX: Cognitive Science Society.
- Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive Psychology*, 72, 54-107.
- Rehder, B. (2018). Beyond markov: accounting for independence violations in causal reasoning. *Cognitive Psychology*, 103, 42-84.
- Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, 50, 264 –314.
- Rehder, B., & Waldmann, M. R. (2017). Failures of explaining away and screening off in described versus experienced causal learning scenarios. *Memory & Cognition*, 45(2), 245-260.
- Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, 140, 109-139.
- Rottman, B. M., & Hastie, R. (2016). Do people reason rationally about causally related events? Markov violations, weak inferences, and failures of explaining away. *Cognitive psychology*, 87, 88-134.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction, and search*. New York: Springer-Verlag.
- Walsh, C. R., & Sloman, S. A. (2004). Revising causal beliefs. In K. Forbus, D. Gentner, & T.

- Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 1423–1429). Mahwah, NJ: Lawrence Erlbaum Associates.
- Walsh, C. R., & Sloman, S. A. (2008). Updating beliefs with causal models: Violations of screening off. In G. H. Bower, M. A. Gluck, J. R. Anderson & S. M. Kosslyn (Eds.), *Memory and mind: A festschrift for Gordon Bower* (pp. 345–358).
- Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language*, 26, 21–52.
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, 136, 82–111.
- Wolff, P., & Barbey, A. K. (2015). Causal reasoning with forces. *Frontiers in Human Neuroscience*, 9.

Appendix

The questionnaire about the bomb explosion problem

An experiment on the impact of bomb exploding on bottles is conducted in an open place. As shown in the figure below, C is a bomb. A, B, D and F represent five bottles which are placed separately. The explosion of bomb C will independently impact each of the five bottles by shockwaves. The bombs have two states. One is the exploding state, which is denoted as “+”, and the other is the unexploded state, which is denoted as “–”. The

bottles have two states, tipping over and vertical state, as shown in the following figures. The probabilities of tipping over of the five bottles given 100 explosions of bomb C are represented below the corresponding bottles respectively.

Please answer the following questions and fill your ratings in brackets. Strength ratings should be given on a scale from 1-7 where larger values indicate greater strength. Probability ratings should be given on a scale from 0% (“impossible”) to 100% (“definitely”). The greater the rating is, the greater the probability of tipping over is.

(1) Given an explosion of bomb C and its effects as shown in the figure below, what is the exploding strength of bomb C? () What is the probability that bottle E will tip over? () What is the probability that bottle F will tip over? ()

(2) Given an explosion of bomb C and its effects as shown in the figure below, what is the exploding strength of bomb C? () What is the probability that bottle E will tip over? () What is the probability that bottle F will tip over? ()

(3) Given an explosion of bomb C and its effects as shown in the figure below, what is the exploding strength of bomb C? () What is the probability that bottle E will tip over? () What is the probability that bottle F will tip over? ()