

Modelling Causal Reasoning under Ambiguity

Yiyun Shou (yiyun.shou@anu.edu.au)

Michael Smithson (michael.smithson@anu.edu.au)

Research School of Psychology, The Australian National University
Canberra, ACT, Australia

Abstract

Causal reasoning under ambiguity requires subjects to estimate and evaluate ambiguous observations. This paper proposes a hierarchical model that accounts for the uncertainty of both the distribution of the functional form selection and distribution of the ambiguity treatment selection. The posterior distribution of the causal estimates is determined by both the functional form and the ambiguity processing strategy adopted by the reasoner. A model is tested in a simulation study for its ability to recover the strategy and functional form adopted by subjects across a range of hypothetical conditions. The model is further applied to the results of an experimental study.

Keywords: Causal reasoning; Ambiguity; Uncertainty

Introduction

A number of models have been proposed in previous several decades to characterize the mechanisms of causal reasoning from covariation information (Lagnado, Waldmann, Hagmayer, & Sloman, 2007). The goals of these models are to account for peoples' inferences about the causal strength between a candidate cause and an effect. These computational and rule-based models assume that subjects observe only unambiguous discrete observations during causal reasoning (e.g., Cheng, 1997; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008; Griffiths & Tenenbaum, 2005).

Problems can arise when the observed causal evidence is ambiguous. In the present paper, we define ambiguous causal evidence as observations where the occurrence of cause and/or effect is unclear; each observation has two possible interpretations: either the cause/effect is present or absent (Pushkarskaya, Liu, Smithson, & Joseph, 2010). Ambiguous observations can result in imprecise probabilistic information which models of rule-based causal reasoning cannot account for.

The way in which the reasoner treats these ambiguous observations could influence the reasoner's judgment on the causal association between the cause and the effect. For example, one reasoner may treat all ambiguous outcomes with the presence of the cause as if they do not have the effect, while another reasoner may treat these outcomes as if they have the effect. The former would have a lower estimate of the causal association than the latter as perceiving there is less evidence supporting that the cause can generate the effect.

Detecting how reasoners treat ambiguous observations, however, can be complicated by the uncertainty in the functional form they adopt to parameterize the causal association. The functional form refers to the way in which the different causes combine in the causal structure (Tenenbaum, Griffiths, & Kemp, 2006; Lucas & Griffiths, 2010). Reasoners hold different assumptions about how the

Table 1: Covariance Information

	Cause Present c^+	Cause Absent c^-
Effect Present e^+	$n(e^+, c^+)$	$n(e^+, c^-)$
Effect Absent e^-	$n(e^-, c^+)$	$n(e^-, c^-)$

alternative causes may interact with the candidate cause in influencing the effect. The assumption of the functional form held by the reasoner determines how subjects will estimate the strength of the causal links between the different causes and the effect.

The goal of this paper is to propose a model of causal reasoning with ambiguous observations, that accounts for the two types of uncertainty: the uncertainty in functional form selection and the uncertainty in the choice of a strategy for processing ambiguous observations. We apply Bayesian inference techniques and the cognitive toolbox approach proposed by Scheibehenne, Rieskamp, and Wagenmakers (2013) to unify the two processes in a single model framework. We first review several computational models of causal reasoning, and explain the connection between these models and functional form learning. We then introduce several possible strategies of ambiguity processing. Subsequently, we propose the model that integrates these two cognitive processes. Next, we perform a simulation study to examine the properties of the model. Finally, we apply the model to empirical data.

Normative Models of Causal reasoning

In the current study, we focus on the elemental causal reasoning, which has only one candidate-cause and one effect (Griffiths & Tenenbaum, 2009). The goal of subjects is to estimate the causal link between the candidate cause and the effect in presence of some background causes. The observations are covariational information outlined in Table 1. We adopt the causal graphical model to illustrate the possible functional forms introduced below. As shown in Figure 1, the background cause (b), candidate cause (c) and the effect (e) are connected by hypothetical causal links. The parameters of the links (w_0 and w_1) indicate the statistical dependencies between b , c and e . Given that there is a link between c and e , three structures can be generated depending on how the three node are combined.

Linear Model The classical model of causal reasoning is ΔP proposed by Jenkins and Ward in 1965. ΔP is defined by $\Delta P = P(e^+|c^+) - P(e^+|c^-)$, where $P(e^+|c^+)$ is the probability of the occurrence of the effect when the cause is present, and $P(e^+|c^-)$ is the probability of the occurrence

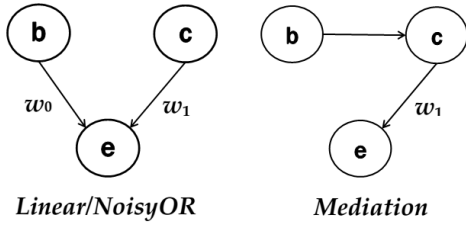


Figure 1: Graphical representation of the functional forms *Notes*. The background cause (b), candidate cause (c) and effect (e) are connected by hypothetical causal links. The links with arrows indicate the dependencies between the nodes.

of the effect when the cause is absent. As illustrated in the first graph in figure 1, ΔP suggests the background cause and the candidate cause linearly combine to generate the effect (Griffiths & Tenenbaum, 2005), :

$$P(e^+|w_0, w_1, b, c) = w_0b + w_1c \quad (1)$$

where b and c are binary indicator variables of the occurrence of the background cause and the candidate cause, w_0 and w_1 are parameters quantifying the strength the links between b and e and between c and e . The causal strength of the candidate cause c is:

$$w_{1,LIN} = P(e^+|c^+) - P(e^+|c^-) \quad (2)$$

Noisy OR: Power PC Model A second way to parameterize the candidate causal structure was logic Noisy OR function proposed by Cheng (1997). A noisy-OR logic function is applied to estimate a generative causal relationship:

$$P(e^+|b, c; w_0, w_1) = w_0b + w_1c - w_0w_1bc \quad (3)$$

The generative causal power or w_1 for a generative causal link is the contribution of the candidate cause independent from the background cause, and is defined by:

$$w_{1,NOR} = \frac{P(e^+|c^+) - P(e^+|c^-)}{1 - P(e^+|c^-)} \quad (4)$$

Single Cause-Full Mediation The last functional form is the one in which subject attributes all effect to the candidate cause: $P(e^+|w_1, c) = w_1c$. One way to interpret this parametrization is that subjects perceive the candidate cause can fully mediate the effect of the background cause, given that the candidate cause occurs after the background cause. The causal strength between the candidate cause and the effect, by fully mediating the effect of other causes, therefore is defined by:

$$w_{1,MED} = P(e^+|c^+) \quad (5)$$

Strategies of Ambiguous Treatment

The two types of reasoners mentioned previously are examples of how two simple strategies may be adopted. We

shall name the two strategies as *positive imputation* (treating all ambiguous observations as if the effect is present) and *negative imputation* (treating all ambiguous observations as if the effect is absent). The possibility that people may treat the incomplete information as negative or positive has also been documented in past research (Garcia-Retamero & Rieskamp, 2008; Lim & Kim, 1992). A third strategy, we name *uniform imputation*, is to treat ambiguous observations as equally likely to have the effect present or absent. This strategy is similar to a "rational" strategy for resolving an interval estimate of a probability by taking the midpoint of the interval.

Mixture Model of Elemental Causal Reasoning

Detecting the strategies adopted by subjects requires distinguishing among the functional forms that subjects may apply in causal reasoning. A relatively low causal rating provided by subjects might be a result of using a linear function in causal reasoning, or applying a negative imputation strategy in treating ambiguous information. Accounting for the functional form could help accurately recover the strategies people use to treat ambiguous information.

We propose a two-step Bayesian hierarchical model to account for both the strategy choice and the function preference. Figure 2 provides a graphical representation. Suppose the i th individual is reasoning about the causal strength y_{ij} of the j th reasoning trial. The individual has observed a_j ambiguous instances, and n_j unambiguous observations that includes the four types of observations: $n_j(e^+, c^+)$, $n_j(e^-, c^+)$, $n_j(e^-, c^-)$ and $n_j(e^+, c^-)$.

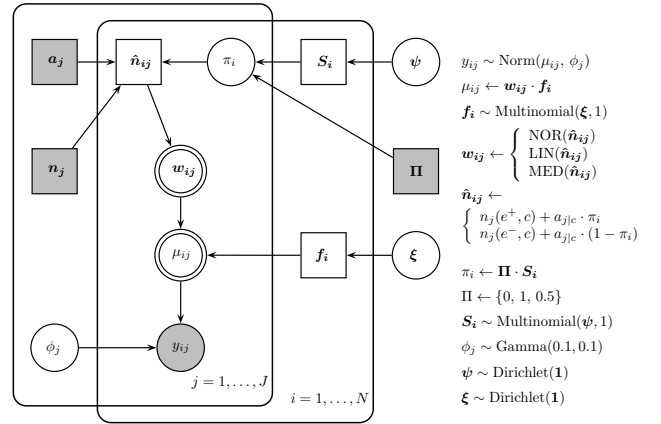


Figure 2: Graphical representation of the model *Note*. The shaded nodes represent observed variables, while the unshaded nodes represent unobserved variables. The square nodes represent discrete variables while the round nodes represent continuous variables. The double-edged nodes represent deterministic variables while the single-edged nodes represent stochastic variables.

Model Stage 1: Processing the Ambiguous Information

Ambiguous observations influence causal judgments by influencing the estimations of the four covariance frequencies, which in turn influences the estimations of the relevant conditional probabilities. The reasoner processes and transforms ambiguous observations, into the forms that can be used for causal judgments

$$\hat{n}_{ij} \begin{cases} \hat{n}_{ij}(e^+, c) = n_j(e^+, c) + a_j \cdot \pi_i \\ \hat{n}_{ij}(e^-, c) = n_j(e^-, c) + a_j \cdot (1 - \pi_i) \end{cases}$$

Where π_i is a strategy selected by the reasoner from the strategy pool Π to impute ambiguous observations. The strategies can be positive imputation ($\pi_1 = 1$), negative imputation ($\pi_2 = 0$) or uniform imputation ($\pi_3 = 0.5$). The choice of the strategy follows $S_i \sim \text{Multinomial}(\psi, 1)$ and ψ is the vector parameter that determines the probabilities of which the strategies would be selected.

Model Stage 2: Causal reasoning The estimated causal strength of each subject μ_{ij} on the j th judgment occasion, is determined by the estimated evidence \hat{n}_{ij} and the preferred functional form f_i . We assume each of subjects adopts one of the three functions outlined previously, $\mu_{ij} = W_{f_i, ij}$. Where $W_{1:3, ij}$ represents the three types of formulation in Equation 1, 2 and 3. The selected function f_i follow $f_i \sim \text{Multinomial}(\xi, 1)$, where ξ is the vector that determines the probabilities for each functional form being adopted.

Simulation Study

We first explored the properties of the model via model recovery studies using simulated data. The simulated subject groups differed in the distributions of the preferred functional forms/ambiguity strategy, and the level of rating variability. We examined to what extent the model could successfully recover the strategy or functional form applied by subjects in various conditions.

In this study, we assumed that subjects apply the same function and the same strategy for ambiguity treatment over time. It is reasonable to assume that subjects are consistent over a short time in an actual experiment. Second, to maximize the distinction between different strategies after being combined with different functional forms, we started from the three simplest strategies.

Simulation Conditions

The simulation study had a factorial design with nine functional form distributions, nine ambiguity strategy distributions, nine levels of variability and two sources of variability. The details of each factor are described below.

Functional form and Ambiguity strategy distributions.

The parameter vector ξ that determines the proportions of the functional forms among the simulated subjects had nine different configurations, including three levels of function dominance (80%, 60% and 40%) by the three types of the functional forms (Linear, NOR and Mediation). For

example, a condition with the dominance level of 80% where the linear function is the dominant function results in $\xi = [0.8, 0.1, 0.1]$. That is, 80% of the subjects adopt the linear function, while 10% of the subjects adopt each of the other non-dominant functions. The proportions of the strategy adoption determined by ψ also had nine configurations, i.e., three levels of dominance crossed with the three strategies.

Between-subject Variability. The first source of response variability is the between-subject variability. There are individual differences in setting up a base-line of responses, and in the ability to reason about causal relationships. We set up the subject-level error term $u_i \sim \text{Normal}(0, \tau^{-1/2})$, where τ is the precision parameter. The lower τ indicates the higher variability of u_i , and the higher dispersion of ratings among subjects. The nine levels of between-subject variability $\tau^{-1/2}$ varied from 0 to 0.2, because the between-functional form variance at different causal covariance levels in the current study ranged from 0.008 to 0.09. To enhance the discrimination between different functional forms, we constrained the between-subject variability to be smaller than the between-function variability.

Within-subject Variability. A second source of variability is the within-subject variability, which indicates the degree of inconsistency of an individual in providing causal judgments. This variability can be regarded as some non-systematic errors when people are judging each trial within the same condition. We define the within-subject error term $\varepsilon_{ij} \sim \text{Normal}(0, \lambda^{-1/2})$. We manipulated the levels of within-individual variability by varying $\lambda^{-1/2}$. The higher $\lambda^{-1/2}$ indicates the higher variability of ε_{ij} , and more inconsistency of subjects in providing judgments. For the same reason as mentioned in previous above, we set up the nine levels of within-subject variability $\lambda^{-1/2}$ to range from 0 to 0.2.

Data generation

For each simulation condition, 20 subjects were simulated, and each of the simulated subject completes 10 causal judgments trials, with different levels of causal strength in the observations. The composition of these observations is shown in Table 2. We chose the covariance levels to ensure (1) the values of $P(e^+|c^+)$ are greater than zero, and are different depending on the strategy of ambiguity treatment, and (2) $P(e^+|c^+)$ will not be smaller than $P(e^+|c^-)$ to keep the causal direction as generative.

First, the estimated contingency cell frequencies were generated by combining the observed evidence and the ambiguity strategy selected by an individual. Next, the aggregated covariance information was passed on to the causal reasoning model component. For the error-free situation, the simulated responses were generated by combining the functional forms (see Equation 2, 4, and 5) and the estimated evidence in each reasoning trial. For conditions

with the between subject variability, $w'_{ij} = w_{ij} + u_i$. For the within-subject variability conditions, $w'_{ij} = w_{ij} + \epsilon_{ij}$.

Estimations and Evaluations

Estimation was performed in JAGS. For each simulated dataset, Markov chain Monte Carlo (MCMC) two-chain processes were run. For each chain, 10000 representative samples were drawn from the posterior distributions, with the first 5000 steps being discarded. The summarized results are based on the 5000 samples, and averaged over the two chains. The summary of the key nodes was obtained. The performance of the model was evaluated by assessing the accuracy of the strategies predicted by the model. Additionally, Bayes factors were calculated to compare the mixture model against a model which assumes all subjects use one single functional form.

Results and Discussion

Figure 3 (upper panel) shows Bayes factors comparing mixture function model against models based on a single function form. For each of the nine functional form distributions and the nine levels of variability, the Bayes factors were averaged over the nine ambiguity distribution conditions. For conditions where a single function form was strongly dominant in the data, the data equally supported the mixture model and the single function form model. The preference for the mixture model over the single functional form model by the Bayes factors increased the preference with the decrease of the dominance of the single function.

Figure 3 (lower panel) displays the accuracies of the model in predicting the ambiguity strategies. For each of the nine ambiguity strategy distributions and the nine levels of variability, the accuracies of ambiguity strategy prediction were averaged over the nine functional form distribution conditions. The proportions of those who correctly recovered the strategy selections were above chance in all conditions. The accuracy of prediction decreased with the increase in the variability in the data.

Experimental Study

We conducted an experimental study and applied the model to predict the strategies used by subjects. We fixed the functional form adopted by subjects by training the subjects to learn a certain type of function form. They were required to apply the given functional form when doing causal reasoning with ambiguous observations.

Method

Participants A total of 77 subjects (42 females, Mean age = 36.44, $SD = 11.2$) were recruited via on-line crowdsourcing platform CrowdFlower. They were paid 80 American cents for their participation. Subjects were randomly assigned to one of the functional form conditions ($N=25$ in the linear function condition, $N=26$ in the NOR function condition, and $N=26$ in the mediation function condition).

Materials Subjects were asked to pretend to be employees in a neurovirology research institution, and their task was to evaluate the effects of a range of chemicals on certain types of viruses which cause neurological diseases. They observed stimuli that indicated the status of the virus and the presence or absence of the chemical, and judged whether the chemical was an activator or inhibitor of the virus.

In each experimental block, subjects were shown two sets of virus samples, each of which had 10 virus samples. The status of a sample virus was either activated or inactivated. They were told that one set had not been exposed to the testing chemical (i.e., the effect status in absence of the cause), and another had been exposed to the testing chemical (i.e., the effect status in presence of the cause).

Their task was to estimate how likely it was that the chemical activated the virus. At the beginning of the task, subjects were instructed according to the functional form condition they were assigned to. Subjects were told about the functional form in the scene, i.e., how the chemical influences the viruses in combination with the background causes. They were instructed how to do causal reasoning¹.

Subjects first went through 10 trials without ambiguous observations. The composition of the covariation frequencies are shown in Table 2. After the 10 trials, they began another 10 trials, in each which one observation in $n_j(e^+, c^+)$ and one observation in $n_j(e^-, c^+)$ were replaced by ambiguous observations, of which the outcome virus was represented by a grayed image with a question mark in the center. The instruction before the ambiguous condition was "In the following several laboratory results, for some reason, the status of some viruses were not clear, they were represented by the grayed picture with a question mark on it as shown in the legend. It does not matter which strategy you select to deal with the unclear viruses, however, it is very essential that you use one strategy consistently through all judgments." At the end of the study, subjects were asked to indicate the strategy that best described their way to process the ambiguous observations.

They were asked to select one of the following choices: (a) regarding the ambiguous virus samples as inactive viruses; (b) regarding the ambiguous virus samples as active viruses; (c) regarding the ambiguous virus samples as half active viruses and half inactive viruses; (d) estimating the ambiguous observations depending on the probability of unambiguous active viruses in presence of the chemical; (e) estimating the ambiguous observations depending on causal link between the unambiguous viruses and the chemical; (f) ignoring the ambiguous viruses; and (g) others.

Results

The proportions of the self-reported strategy among subjects were 45.46% for the negative imputation ($N = 35$), 19.48% for the positive imputation ($N = 15$), 22.08% for the uniform

¹The instructions of the rules are available in the online supplemental materials at <http://goo.gl/PdJWMy>

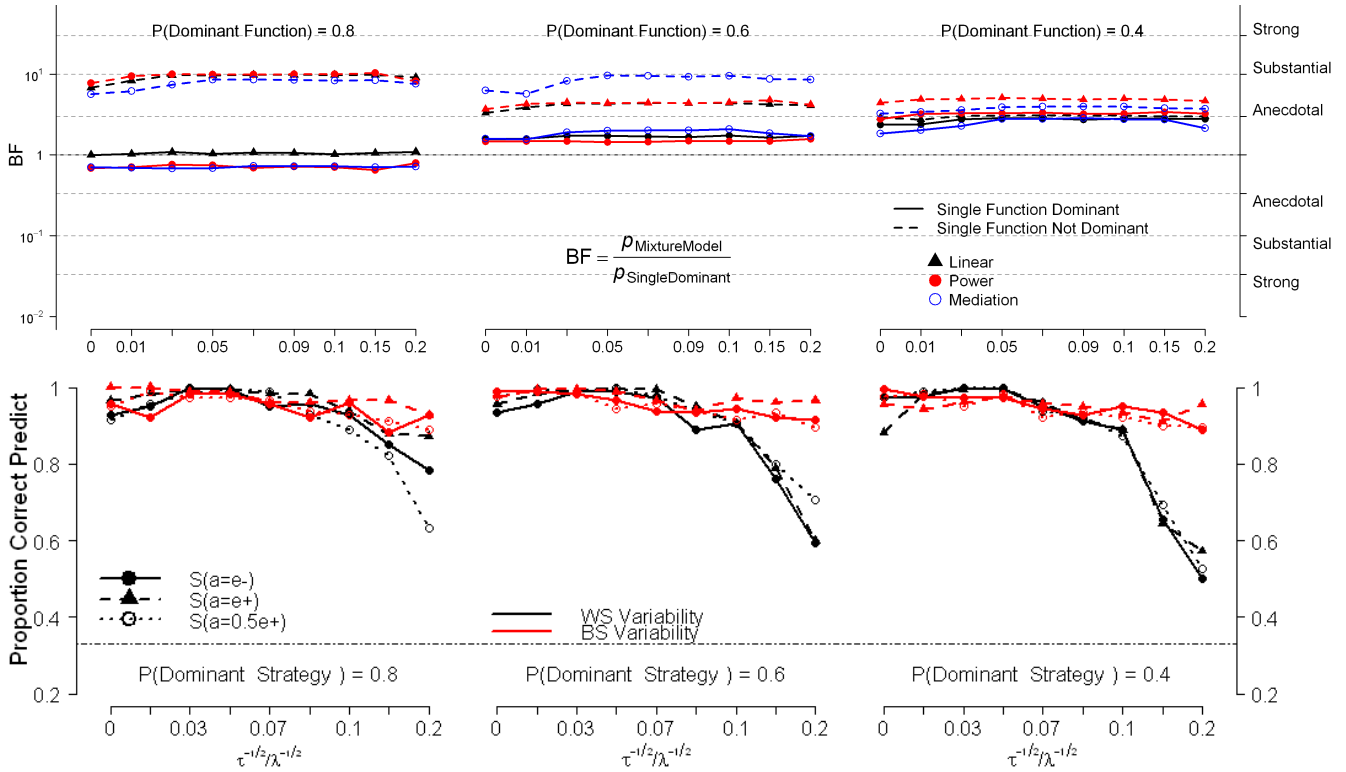


Figure 3: Bayes Factors for Comparing the Mixture Model against Single Function Models(upper panel) and Accuracy of the Model in Predicting Ambiguity strategies across Simulation Conditions(lower panel)

Note. The annotated proportion indicates the level of dominance of a strategy/functional form adopted in the data. The dashed horizontal line in the lower panel is the chance level (0.33).

Table 2: Covariance information in the experiment

	$n_j(e^+, c^+)$	$n_j(e^-, c^+)$	$n_j(e^+, c^-)$	$n_j(e^-, c^-)$
C1	5	5	2	8
C2	7	3	5	5
C3	7	3	2	8
C4	8	2	3	7
C5	3	7	2	8
C6	6	4	4	6
C7	9	1	6	4
C8	3	7	1	9
C9	6	4	2	8
C10	9	1	3	7

imputation ($N = 17$), and 12.99% for all other strategies ($N=2$ for applying the unambiguous observations, and $N=7$ for others). We implemented the mixture model to estimate the functional forms and ambiguity strategy simultaneously. The functional forms predicted by the model matched all those assigned to subjects. The ambiguity strategies predicted by the model agreed with 72.7% of the subjects' self-report strategies, Cohens kappa = 0.60 (83.58 % of the 67 subjects who reported the first three strategies, Cohens kappa = 0.74).

We then fit the model to the entire data. The estimation results of the ξ and ψ were displayed in Table 3. The estimated proportion of the strategy selection generally matched the self-reported strategy. We compared the mixture model against each of the single strategy model with controlling the functional forms. Bayes factors were 2.3 for the mixture model over the negative-imputation-only model, 9.9 over the positive-imputation-only model, and 3.35 over the uniform-imputation-only model, respectively.

Table 3: Parameter Estimation for Experimental Data

	Model	SE	2.5%	97.5%	Self-Report
ξ_{Linear}	0.32	0.05	0.22	0.43	0.32
ξ_{NOR}	0.34	0.05	0.24	0.44	0.34
ξ_{Med}	0.34	0.05	0.24	0.45	0.34
$\psi_{Negative}$	0.48	0.06	0.37	0.59	0.45
$\psi_{Positive}$	0.16	0.04	0.09	0.25	0.19
$\psi_{Uniform}$	0.36	0.05	0.26	0.47	0.22

Note. The self-reported of the strategy selection was out of the entire data including subjects who reported strategies other than the three main strategies.

Discussion

Treating ambiguous observations as the negative cues (i.e., the effect was absent) was most preferred by the subjects, followed by the uniform imputation strategy and positive imputation. The results are in line with some studies of missing data treatment in which people were more likely to treat the missing information as negative cues in decision making (Lim & Kim, 1992). The findings can also be related to studies of ambiguity aversion, where people generally disprefer a gamble when the probability is ambiguous, and are pessimistic about the probability of a reward from such a gamble. This suggests that processing ambiguous observations in causal reasoning may share mechanisms with processing ambiguous and missing information in decision making.

General Discussion and Future Extension

We applied Bayesian hierarchical modelling approach to account for individual differences in applying functional forms to reasoning about causal relationships and adopting strategies for processing ambiguous observations. This current model can distinguish a case where a subject applies the mediation function and is pessimistic about the outcome probability of the ambiguous observations, from a case where a subject applies a linear function and is optimistic about the outcome probability of the ambiguous observations. The risk of misidentifying strategy of ambiguity processing can be reduced.

The current model can be readily extended in several ways. First, the collection of the strategies can be expanded to include more strategies. The three strategies illustrated in this paper do not require the integration of the observed evidence when treating the ambiguous information, and are independent from the unambiguous observations. Other possible strategies, as noted by Garcia-Retamero and Rieskamp (2008), may include ignoring ambiguous observations and averaging the unambiguous observations to estimate the probability of the effect occurring in ambiguous observations.

Second, the model can incorporate more functional forms, such as Noisy-AND-NOT function in reasoning the preventive causal link (Lu et al., 2008). The inclusion of functional forms especially for preventive causal reasoning may also need to take into account interactions between the functional form and strategy selection.

Furthermore, we assume subjects applied the same strategy or functional form through the time of course. Future study may consider the possibility that subjects may be inconsistent in applying for strategies and investigate the factors that may contribute to the inconsistency.

Finally, a parameter that represents how subjects may weigh ambiguous observations may be considered. It is possible that subjects perceive the ambiguous observations as weak evidence in comparison to the unambiguous observations, especially when the proportion of observations

that are ambiguous is small.

The current model provides a framework to explain causal reasoning with ambiguous observations. The model can be applied to improve understanding of how different factors such as the available cognitive resources, distribution of ambiguous observations, and prior knowledge about the reasoning scenario influence subjects' treatment of ambiguous observations.

References

- Cheng, P. W. (1997). From Covariation to Causation: A Causal Power Theory. *Psychological review*, *104*(2), 367–405.
- Garcia-Retamero, R., & Rieskamp, J. (2008). Adaptive mechanisms for treating missing information: A simulation study. *The Psychological Record*, *58*, 547–568.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive psychology*, *51*(4), 334–384.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological review*, *116*(4), 661–716.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond Covariation. In *Causal learning: Psychology, philosophy, and computation* (p. 154).
- Lim, J.-S., & Kim, J. (1992). Impact of consumers' confidence in judgements about missing information on product evaluations. *Journal of Business Research*, *25*(3), 215–229.
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological review*, *115*(4), 955–984.
- Lucas, C. G., & Griffiths, T. L. (2010). Learning the Form of Causal Relationships Using Hierarchical Bayesian Models. *Cognitive Science*, *34*(1), 113–147.
- Pushkarskaya, H., Liu, X., Smithson, M., & Joseph, J. E. (2010). Beyond risk and ambiguity: deciding under ignorance. *Cognitive, affective & behavioral neuroscience*, *10*, 382–391. doi: 10.3758/CABN.10.3.382
- Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: a Bayesian hierarchical approach. *Psychological review*, *120*, 39–64.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318.