

Sequential diagnostic reasoning with independent causes

Marko Tešić (mtesic02@mail.bbk.ac.uk) & Ulrike Hahn (u.hahn@bbk.ac.uk)

Department of Psychological Sciences
Birkbeck, University of London
Malet Street, London, WC1E 7HX, UK

Abstract

In real world contexts of reasoning about evidence, that evidence frequently arrives sequentially. Moreover, we often cannot anticipate in advance what kinds of evidence we will eventually encounter. This raises the question of what we do to our existing models when we encounter new variables to consider. The standard normative framework for probabilistic reasoning yields the same ultimate outcome whether multiple pieces of evidence are acquired in sequence or all at once, and it is insensitive to the order in which that evidence is acquired. This equivalence, however, holds only if all potential evidence is incorporated in a single model from the outset. Hence little is known about what happens when evidence sets are expanded incrementally. Here, we examine this contrast formally and report the results of the first study, to date, that examines how people navigate such expansions.

Keywords: sequential diagnostic reasoning; sequential causal structure learning; causal Bayesian networks; order effects.

Introduction

Tom wakes up one morning and notices a rash on his skin. He does not think the rash is a big deal, but after a couple of days the rash is still present so he decides to see a doctor. Before he visits a doctor he thinks that the rash is either caused by a bacterial or a viral infection or, perhaps, both. The doctor agrees with him that the rash could be caused by a bacterial and/or a viral infection. However, she additionally informs Tom that he also has a swelling he didn't notice, which can also be caused by a bacterial and/or a viral infection. Furthermore, she tells him that either type of infection is more likely to cause the swelling than the rash. How do (should) Tom and the doctor revise their beliefs about multiple independent causes given multiple pieces of evidence of different diagnosticity?

From a normative standpoint, many would argue that the answer is encoded in the causal Bayesian networks (CBNs): directed acyclic graphs with nodes representing variables (causes and effects) and arrows representing probabilistic and causal relations between the nodes (Pearl, 2009, 1988; Neapolitan, 2003). Here one would build a 4-node CBN with 2 common effects and 2 independent causes.¹ For instance, the CBN in Figure 1 would model the situation we described above: C_1 = viral infection, C_2 = bacterial infection, E_1 = rash, and E_2 = swelling.

To fully parameterize CBN from Figure 1, one needs to specify the following probabilities:

$$P_1(C_1) = c_1 \quad , \quad P_1(C_2) = c_2$$

$$P_1(E_1 | C_1, C_2) = \alpha_1 \quad , \quad P_1(E_1 | C_1, -C_2) = \beta_1$$




Figure 1: CBN with 2 independent causes and 2 common effects.

$$P_1(E_1 | -C_1, C_2) = \gamma_1 \quad , \quad P_1(E_1 | -C_1, -C_2) = \delta_1 \quad (1)$$

$$P_1(E_2 | C_1, C_2) = \alpha_2 \quad , \quad P_1(E_2 | C_1, -C_2) = \beta_2$$

$$P_1(E_2 | -C_1, C_2) = \gamma_2 \quad , \quad P_1(E_2 | -C_1, -C_2) = \delta_2$$

$P_1(C_1)$ and $P_1(C_2)$ are usually referred to as the prior probability of the two causes and the remaining probabilities as being part of the conditional probabilities tables (CPTs) for the two effects. The doctor then could use this CBN to update her beliefs about the probability that Tom has a viral infection after learning that Tom has a rash by calculating $P_1(C_1 | E_1)$. After additionally learning that Tom also has swelling the doctor could further update her probability of Tom having a viral infection by calculating $P_1(C_1 | E_1, E_2)$ (similarly for the bacterial infection).

However, it is somewhat accidental that Tom first noticed the rash and not the swelling. He could have plausibly first seen the swelling and gone to the doctor and then noticed the rash. Would the CBN calculation be different in this scenario? It depends. If the rash and the swelling are not equally diagnostic of the two causes as is suggested by the example, then it is possible that $P_1(C_1 | E_1) \neq P_1(C_1 | E_2)$, in which case the doctor's degrees of belief about a viral infection after first learning that Tom has swelling would not be equal to those where she first learned about the rash. However, after learning the second effect the order in which the effect appear no longer matters: that is, $P_1(C_1 | E_1, E_2)$ is always equal to $P_1(C_1 | E_2, E_1)$.

It is then empirically interesting to investigate whether people are sensitive to these different orders of effects and whether they update the causes differently depending on the order in which the effects appear. Studies on sequential di-

¹Hayes, Hawkins, Newell, Pasqualino, and Rehder (2014) have used a dynamic CBN to model these kinds of situations. However, in this paper we employ static CBNs as there are no significant differences in the formalism in this case.

agnostic reasoning have sought to tackle exactly these issues (see Meder & Mayrhofer, 2017b; Hogarth & Einhorn, 1992). They presented participants with a sequence of effects and asked them to reason from multiple effects to causes either with each effect they learned (step-by-step procedure) or after they learned about the whole sequence of effects (end-of-sequence procedure) (see Hogarth & Einhorn, 1992; Rebitschek, Bocklisch, Scholz, Krems, & Jahn, 2015). Their studies were primarily interested in investigating primacy effects (most of the evidential weight is given to the first piece of evidence) and recency effects (most of the evidential weight is given to the most recent pieces of evidence). Meder and Mayrhofer (2017a) investigated sequential diagnostic reasoning by providing participants with verbal information regarding the strengths of the causes instead of a more quantitative information (like the CPTs) and found that participants are remarkably accurate in their judgements. However, all these studies investigated only situations where the causes were mutually exclusive and exhaustive causes (which would be modeled as one node for all causes). Hayes et al. (2014) investigated a scenario where two symptoms could be produced by two independent causes. However, in their study both effects had exactly the same diagnosticity (i.e. the same CPT) and for that reason there are no order effects, i.e. it does not matter whether we learn first E_1 or E_2 , $P_1(C_1 | E_1) = P_1(C_1 | E_2)$.

One of the goals of this paper is to empirically investigate people's ability to reason diagnostically from multiple effects with different diagnosticities (CPTs) to multiple independent causes. More specifically, we aim to test how people's judgements compare to the normative answer from CBNs such as the one in Figure 1 by manipulating the way in which multiple pieces of the evidence of different diagnosticity are presented (in a particular order or at the same time) and the way judgements about the causes are elicited from the participants (step-by-step (SbS) or all-at-once (AaO)).

Another interesting issue emerges when reasoning with independent causes. Not only can we learn the evidence sequentially, but we can sequentially learn about new variables that may influence our beliefs about the causes. In technical parlance, we may need to expand the algebra. Consider Tom from our example. Initially Tom only knew about his rash and, based on that knowledge, he updated his probabilities of the two causes. Unlike the doctor, Tom did not even know that the two types of infection could also cause swelling. It is only after he visited his doctor that he learned about the another potential effect and the occurrence of that effect. At the time he only knew about the rash he updated the probabilities of the two causes on the basis of a CBN model with only three nodes: two independent causes and one common effect while the doctor always had in mind the CBN from Figure 1. Despite operating with two different CBNs, both Tom and the doctor would arrive at the same probabilities (assuming the same priors and CPTs for the effect) at this first step. The next step is, however, crucial. After learning about the

swelling, the doctor would simply learn the new piece of evidence and update the probabilities of the causes based on the CBNs from Figure 1. Tom, by contrast, might do one of two things: (1) forget about his original 3-node network and create a new 4-node one like the one in Figure 1 in which case he would arrive at the same estimates as the doctor; or (2) take his (and doctors) previous estimates of the two causes based on only one piece of evidence and take them as new priors in his new 3-node network with the second piece of evidence as a common effect (see Figure 2). In the latter case he would be 'splitting' the CBNs from Figure 1 into two 3-node networks.




Figure 2: 'Split' CBN from E_1 to E_2

$$\begin{aligned}
 P_1(C_1) &= e_1 & , & & P_1(C_2) &= e_2 \\
 P_2(C_1^*) &= P_1(C_1 | E_1) & , & & P_2(C_2^*) &= P_1(C_2 | E_1) \\
 P_1(E_1 | C_1, C_2) &= \alpha_1 & , & & P_1(E_1 | C_1, -C_2) &= \beta_1 \\
 P_1(E_1 | -C_1, C_2) &= \gamma_1 & , & & P_1(E_1 | -C_1, -C_2) &= \delta_1 \quad (2) \\
 P_2(E_2 | C_1^*, C_2^*) &= \alpha_2 & , & & P_2(E_2 | C_1^*, -C_2^*) &= \beta_2 \\
 P_2(E_2 | -C_1^*, C_2^*) &= \gamma_2 & , & & P_2(E_2 | -C_1^*, -C_2^*) &= \delta_2
 \end{aligned}$$

Eq. (2) specify the priors and the CPTs of the two networks. Although one might intuitively think that Tom will arrive at the same probabilities as the doctor even in the case where he models the situation as in Figure 2, that turns out to be true only under very specific conditions, some of which may violate common assumptions in causal Bayesian reasoning (see Appendix A). Less technically, this is because once one learns evidence (E_1) and updates the probabilities of the two causes (C_1 and C_2) in a common-effect CBN, the two previously independent causes become dependent: although $P_1(C_1 | C_2) = P_1(C_1)$, generally $P_1(C_1 | C_2, E_1) \neq P_1(C_1 | E_1)$. This dependency is preserved in the full CBN network in Figure 1 even *before* one learns the second piece of evidence (E_2) and again updates the probabilities of the two causes. However, in the lower 3-node CBN in Figure 2 the dependency is lost since it is assumed that C_1^* and C_2^* are independent *before* observing E_2 . Therefore, the final probability estimates of the two causes, i.e. their estimates *after* learning both pieces of evidence, will most likely diverge on the two different modeling strategies. More specifically, the final estimates of the two causes will always be higher according to the 'split' CBN in Figure 2 than those according to the full one in Figure 1 precisely because the full one accounts for the above-mentioned dependency and the 'split' one does not. Moreover, when the diagnosticity of the two pieces of evidence is different (as is the case in this study), the height of the final estimates in the 'split' CBN will depend on the order the evidence is observed: learning E_1 then E_2 will result

in the final estimates of the causes that are different to those that result from learning E_2 then E_1 (as previously mentioned, whether we learn E_1 first then E_2 or vice versa does not affect the final probability estimates of the causes in the full CBN). It is also worth pointing out that this divergence only happens when the causes are independent. If the causes are mutually exclusive and exhaustive, one can safely ‘split’ the full network into multiple ones without worrying about ending up with different estimates (see Appendix B).

To the best of our knowledge no study has yet investigated sequential diagnostic reasoning with sequentially learning the algebra. In the literature mentioned above participants were presented with all the variables and the causal/probabilistic information related to them before they started making judgements about the causes. Even in such contexts, it is worth looking at order effects because it has long been recognized that order effects may be particularly diagnostic with respect to the processes underlying the formation of a judgment. Specifically, there is a long literature concerned with order effects in contexts such as impression formation (Anderson, 1965) or numerical estimation (Jacowitz & Kahneman, 1995). However, our concerns in this paper go beyond this. We are interested in examining how reasoners fare in probabilistic reasoning contexts where they are faced with entirely new variables. This issue has, to the best of our knowledge, not been explored. In many scientific and everyday situations we must make judgements about potential causes given effects without being aware of other potential effects that could also inform our judgements. The main aim of this study was to examine how people reason with multiple pieces of evidence when they successively learn not just that some piece of evidence obtains, but also that there is another potential piece of evidence not known before. We compared participants’ estimates to both the full network’s predictions (Figure 1) and the ‘split’ networks’ predictions (Figure 2).

Experiment overview

In the present experiment we investigated the influence of manipulating algebra and evidence learning on probabilistic judgements of the two independent causes. Participants were prompted to reason with either the full 4-node model (Figure 1) from the outset or they learned in stages that there is another possible effect of the two causes. Further, participants either observed the effects in one of the two sequences or they observed both effects at once. The prior probabilities of the cases and CPTs of the effects were the same in all conditions: $P(C_1) = P(C_2) = 0.15$; $P(E_1 | C_1, C_2) = 0.99$, $P(E_1 | C_1, \neg C_2) = P(E_1 | \neg C_1, C_2) = 0.7$, $P(E_1 | \neg C_1, \neg C_2) = 0$; $P(E_2 | C_1, C_2) = 0.6$, $P(E_2 | C_1, \neg C_2) = P(E_2 | \neg C_1, C_2) = 0.2$, $P(E_2 | \neg C_1, \neg C_2) = 0$. For simplicity the priors of the causes were the same and the CPTs of the effects reflected different diagnosticities of the two effects.

Methods

Participants and design

A total of 271 participants ($N_{\text{MALE}} = 101$, $M_{\text{AGE}} = 32.1$ years; one participant identified as neither male nor female) were recruited from Prolific Academic (www.prolific.ac). All participants were native English speakers who gave informed consent and were paid £1.25 for partaking in the present study, which took on average 13.9 minutes to complete. Participants were randomly assigned to one of the 2 (algebra: full or sequential) \times 3 (evidence learning: all-at-once (AaO), step-by-step from E_1 to E_2 (SbS1), or step-by-step from E_2 to E_1 (SbS2)) = 6 between-participants groups (one group with 44 participants, 3 groups with 45 participants, and 2 groups with 46 participants).

Materials

All participants were given the same cover story wherein rain (C_1) and a lawn sprinkler (C_2) (two binary and independent variables) could cause a wet lawn (E_1) and/or a wet exterior house wall (E_2) (a version of the cover story can be found in Pearl, 1988). The participants in AaO condition completed an online inference questionnaire comprising of 10 comprehension questions (2 about the priors of the causes and 8 about the CPTs) and 2 test questions (one relating to $P(C_1 | E_1, E_2)$ and one to $P(C_2 | E_1, E_2)$). Everyone else completed the same 10 comprehension questions and 4 test questions (relating to $P(C_1 | E_i)$, $P(C_2 | E_i)$, $P(C_1 | E_i, E_j)$, and $P(C_2 | E_i, E_j)$).

Procedure

In the full algebra condition, the participants were initially presented with a causal cover story (both in a textual and a visual form) which explained the relations between variables and probabilistic information relating to the priors of both causes (priors were textually communicated as a percentage chance). They were then asked 2 priors comprehension questions. Following that, participants were told the CPTs of the two pieces of evidence (also textually communicated as a percentage chance) and subsequently asked 8 comprehension questions regarding the CPTs (in a random order). After completing the comprehension questions, participants in the AaO condition learned that both pieces of evidence occurred and were prompted to answer 2 test questions (one for each cause) presented in the same order. Participants in the SbS conditions first learned about one piece of evidence and answered 2 test questions relating to the 2 causes and then learned that the second piece of evidence occurred and asked final 2 questions. When answering the test questions participants were reminded of the priors of the causes and the CPTs of each piece of evidence, as well as their previous estimates of the two causes (in the SbS conditions).

Participants in the sequential algebra condition were initially told a cover story (both in a textual and a visual form) including only two causes and one effect. As in the full algebra condition, they were told the priors of the causes (percentage chance) and asked 2 priors comprehension questions.

In contrast to the full algebra contention, they were then told CPTs (percentage chance) regarding only one piece of evidence and completed 4 comprehension questions related to CPTs (in both the AaO and the SbS conditions). This was followed by 2 test questions relating to the probability of the causes given that one piece of evidence was observed (only in the SbS conditions). Participants then additionally learned that there is another piece of evidence potentially relevant to the probability estimates of the two causes. They learned the CPTs for the second piece of evidence and completed 4 comprehension questions followed by 2 test questions prompting them to estimate their confidence in the causes happening given the additional piece of evidence obtained. Again, participants were reminded of the priors of the causes, CPTs (but only for the current piece of evidence), and their previous estimates of the two causes (in the SbS conditions). In the AaO, after completing the first 4 comprehension questions participants were not told that the evidence obtained. Rather, they went on to learn that there is another potentially relevant piece of evidence, completed additional 4 comprehension questions, and subsequently told that both pieces of evidence obtained. After that, participants were reminded of the priors, CPTs (for the both pieces of evidence) and completed 2 test questions regarding the probabilities of the two causes.

In all conditions the test questions prompted participants to provide percentage confidence (0–100%) of C_i given one or two effects. For example, after learning that E_1 occurred, they were asked (in SbS1 condition) a diagnostic reasoning questions: “How confident are you that it **rained** overnight now that you know that the lawn is wet?” After additionally learning E_2 occurred they were asked: “How confident are you that it **rained** overnight now that you know that both the lawn and the house wall are wet?” (the full algebra condition) or “How confident are you that it **rained** overnight now that you know that the house wall is also wet?” (the sequential algebra condition). All participants provided explanations for each answer to the test questions.

Results

All the participants’ responses to the test questions are plotted in Figure 3. To test the effect of the algebra and the evidence learning conditions on participants estimates on the test questions, we built a linear mixed effects model using the lme4 package (Bates, Mächler, Bolker, & Walker, 2014). The model had two fixed effects, Algebra and Evidence learning, with a random intercept for each participant (there was no random slope for participant since algebra and evidence learning conditions vary between participants). We found a main effect of Evidence learning but no main effect of Algebra (see Table 1). We also found no interaction between Algebra or Evidence learning. However, likelihood ratio tests showed that including the predictors in the model does not improve model fit compared to just having an intercept as a predictor ($\chi^2(3) = 6.11, p = 0.11$). That is, the data grand

mean fits the data no worse than the model which includes both predictors.

Table 1: Linear mixed effect model results

A=Algebra; EL=Evidence learning

	Estimate	95% CI	t-value	p
A	-6.51	[-17.76, 4.73]	-1.13	0.26
EL	-0.53	[-1.03, -0.03]	-2.1	0.04*
A × EL	3.28	[-17.76, 4.73]	1.29	0.2

A finer grained analyses on the data within each group showed a significant difference between $P(C_1 | E_i)$ and $P(C_1 | E_i, E_j)$ in the full algebra SbS1 condition ($t(44) = -4.04, p = 0.0002$); in the full algebra SbS2 condition both between $P(C_1 | E_i)$ and $P(C_1 | E_i, E_j)$ ($t(45) = -4.87, p < 0.0001$) and $P(C_2 | E_i)$ and $P(C_2 | E_i, E_j)$ ($t(45) = -2.98, p = 0.005$); as well as in the sequential algebra SbS2 condition between $P(C_1 | E_i)$ and $P(C_1 | E_i, E_j)$ ($t(45) = -5.57, p < 0.0001$) and between $P(C_2 | E_i)$ and $P(C_2 | E_i, E_j)$ ($t(45) = -6.13, p < 0.0001$). No significant differences in the sequential SbS1 condition.

Further analyses showed that none of the $P(C_1 | E_i, E_j)$ and $P(C_2 | E_i, E_j)$ are significantly different across the levels of the evidential learning condition whereas some $P(C_2 | E_i)$ are: in the full algebra condition $P(C_2 | E_i)$ in SbS2 and SbS1 are statistically different, $t(89) = -2.09, p = 0.04$, as well as $P(C_2 | E_i)$ in the sequential algebra condition SbS2 and SbS1 $t(88.5) = -2.51, p = 0.014$, with those in SbS1 having higher means. Combining these results from those above regarding participants estimates withing each group suggests that (i) people are sensitive to the different orders the pieces of evidence of different diagansticity were presented and (ii) that their estimates go against both the full CBN and the ‘split’ CBNs (qualitative) predictions since the differences $P(C_1 | E_i, E_j) - P(C_1 | E_i)$ and $P(C_2 | E_i, E_j) - P(C_2 | E_i)$ are larger in SbS2 condition than in SbS1 condition whereas the full CBN and the ‘split’ CBN predict exactly the opposite (see Figure 3).

A closer look at the data distributions in Figure 3 reveals the driving force of the results; namely, that participants’ responses are highly clustered. Three clustering points (‘20%’, ‘60%’, and ‘70%’) seem to correspond to the probability values one finds in the CPTs for the effects. One clustering point corresponds to the priors of the causes (‘15%’). The largest clustering point seems to be around the ‘50%’ mark. Table 3 shows a frequency of responses around ($\pm 2\%$) the clustering points. The data captured in Table 3 amounts to $\approx 67\%$ of all data.

Finally, to assess the fit of each model to the data, we calculated mean squared errors (MSEs) for each model across the two algebra conditions.² Given the above-mentioned cluster-

²Note that the ‘split’ CBN does not have a unique prediction for AaO condition (see Figure 3). In calculating the MSE for that model we included the prediction that has the lower MSE.

ing around particularly the ‘50%’ mark, we additionally calculated the MSEs for a simple model that included the correct priors (same as in both the full CBN and the ‘split’ CBN modeling), but has 50% as a response to all test questions. The results are presented in Table 2.

Table 2: MSEs for the full CBN, ‘split’ CBN, and ‘50%’ model in the full and sequential algebra conditions

	Full algebra	Sequential algebra
Full CBN	621.18	536.94
‘Split’ CBN	778.93	701.97
‘50%’ model	573.73	496.65

The best fitting model of the three was the simple ‘50%’ model, further confirming the clustering effect around the ‘50%’ mark and the results of the linear mixed effect model. The full CBN model was a better fit than the ‘split’ CBN model of both the full algebra condition data and sequential algebra condition data. All three models fit better the sequential algebra condition data than the full algebra condition data suggesting a difference between the two conditions. However, according to the linear mixed effect model that difference is not statistically significant.

Discussion

The general goal of the paper was twofold. First, we sought to explore new avenues in sequential diagnostic reasoning by investigating peoples causal judgements with multiple independent causes and multiple pieces of evidence of different diagnosticity. To this effect we found that people are sensitive to the order of presentation of the different pieces of evidence. However, although there was a trend in increasing the probabilities of the causes after finding out that the second piece of evidence obtained (in accordance with both the full and the ‘split’ CBN model), the (qualitative) predictions of both models regarding the amount of increase in each order go against the participants’ mean estimates.

Second, we introduced the issue of the novel variables in sequential reasoning and the practical challenge it presents. In the first empirical study on this issue, we found that people update almost identically when they are presented with the full algebra and when the algebra is expanded sequentially. In principle, this lack of difference could mean either that people are very good at this expansion, or that they inappropriately treat the full model in a sequential, local fashion. The MSE analysis showed that the full CBN model is a better fit than the ‘split’ CBN model across board supporting the latter option. However, the significant clustering in our data and the fact that the ‘50%’ model fitted the data better than either the full or the ‘split’ CBN model suggest that participants employed different strategies in answering our test questions. Some of these seem indicative of well-established errors in human causal/probabilistic reasoning such as ‘the inversion fallacy’ where people confuse $P(A | \neg B)$ with $P(B | A)$ (Nance & Morris, 2002) or more recently identified errors such as ‘the

zero-sum fallacy’ where people treat evidence as a zero-sum game in which alternative independent hypotheses compete for evidential support which may lead to splitting the probability space between the hypotheses (Pilditch, Fenton, & Lagnado, 2019). The prevalence of such errors may mask other differences that would emerge across those contexts. In particular, systematic differences may yet be found in more naturalistic scenarios where there are no explicit numbers for people to hold on to. This should be pursued in future work.

Acknowledgments

This research has been partly supported by the Humboldt Foundation’s “Anneliese Meier Research Award” to Ulrike Hahn.

References

- Anderson, N. H. (1965). Primacy effects in personality impression formation using a generalized order effect paradigm. *Journal of personality and social psychology*, 2(1), 1–9.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.
- Hayes, B. K., Hawkins, G. E., Newell, B. R., Pasqualino, M., & Rehder, B. (2014). The role of causal models in multiple judgments under uncertainty. *Cognition*, 133(3), 611–620.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24(1), 1–55.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21(11), 1161–1166.
- Meder, B., & Mayrhofer, R. (2017a). Diagnostic causal reasoning with verbal information. *Cognitive psychology*, 96, 54–84.
- Meder, B., & Mayrhofer, R. (2017b). Diagnostic reasoning. In M. R. Waldmann (Ed.), *Oxford handbook of causal reasoning* (pp. 433–458). Oxford University Press New York.
- Nance, D. A., & Morris, S. B. (2002). An empirical assessment of presentation formats for trace evidence with a relatively large and quantifiable random match probability. *Jurimetrics*, 42, 403–448.
- Neapolitan, R. E. (2003). *Learning bayesian networks*. Pearson Prentice Hall Upper Saddle River, NJ.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA: Morgan Kaufman.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Pilditch, T. D., Fenton, N., & Lagnado, D. (2019). The zero-sum fallacy in evidence evaluation. *Psychological Science*. Retrieved from <https://doi.org/10.1177/0956797618818484>
- Rebitschek, F. G., Bocklisch, F., Scholz, A., Krems, J. F., & Jahn, G. (2015). Biased processing of ambiguous symptoms favors the initially leading hypothesis in sequential diagnostic reasoning. *Experimental psychology*, 62(5), 287–305.




Figure 3: Distributions of participants' responses to the test questions. Error bars are 95% confidence intervals.

Appendix A

We adopt the following convention: $\bar{a} = 1 - a$.

Theorem 1. $P_1(C_1 | E_1, E_2) = P_2(C_1^* | E_2)$ if and only if (i) $\alpha_1 \delta_1 = \beta_1 \gamma_1$ or (ii) $\alpha_2 = \beta_2$ and $\gamma_2 = \delta_2$.

Proof.

$$\begin{aligned} P_1(C_1 | E_1, E_2) &= \frac{P_1(C_1, E_1, E_2)}{P_1(E_1, E_2)} \\ &= \frac{P_1(C_1) \sum_{C_2} P_1(E_1 | C_1, C_2) P_1(E_2 | C_1, C_2) P_1(C_2)}{\sum_{C_1, C_2} P_1(E_1 | C_1, C_2) P_1(E_2 | C_1, C_2) P_1(C_1) P_1(C_2)} \\ &= \frac{A_1}{A_1 + A_2} \end{aligned}$$

$$A_1 := c_1 (\alpha_2 \alpha_1 c_2 + \beta_2 \beta_1 \bar{c}_2)$$

$$A_2 := \bar{c}_1 (\gamma_2 \gamma_1 c_2 + \delta_2 \delta_1 \bar{c}_2)$$

$$\begin{aligned} P_2(C_1^* | E_2) &= \frac{P_2(C_1^*, E_2)}{P_2(E_2)} \\ &= \frac{P_2(C_1^*) \sum_{C_2} P_2(E_2 | C_1^*, C_2^*) P_1(C_2^*)}{\sum_{C_1^*, C_2^*} P_2(E_2 | C_1^*, C_2^*) P_2(C_1^*) P_2(C_2^*)} \\ &= \frac{P_1(C_1 | E_1) \sum_{C_2} P_1(E_2 | C_1, C_2) P_1(C_2 | E_1)}{\sum_{C_1, C_2} P_1(E_2 | C_1, C_2) P_1(C_1 | E_1) P_1(C_2 | E_1)} \\ &= \frac{B_1}{B_1 + B_2} \end{aligned}$$

$$B_1 := c_1 (\alpha_1 c_2 + \beta_1 \bar{c}_2) \cdot$$

$$\cdot [\alpha_2 c_2 (\alpha_1 c_1 + \gamma_1 \bar{c}_1) + \beta_2 \bar{c}_2 (\beta_1 c_1 + \delta_1 \bar{c}_1)]$$

$$B_2 := \bar{c}_1 (\gamma_1 c_2 + \delta_1 \bar{c}_2) \cdot$$

$$\cdot [\gamma_2 c_2 (\alpha_1 c_1 + \gamma_1 \bar{c}_1) + \delta_2 \bar{c}_2 (\beta_1 c_1 + \delta_1 \bar{c}_1)]$$

Let $\Delta_1 := P_1(C_1 | E_1, E_2) - P_2(C_1^* | E_2)$. Then

$$\begin{aligned} \Delta_1 &= \frac{A_1 (B_1 + B_2) - B_1 (A_1 + A_2)}{(A_1 + A_2) (B_1 + B_2)} \\ &= \frac{A_1 B_1 + A_1 B_2 - A_1 B_1 - A_2 B_1}{P_1(E_1, E_2) P_2(E_2)} = \frac{A_1 B_2 - A_2 B_1}{P_1(E_1, E_2) P_2(E_2)} \\ &= \frac{c_1 \bar{c}_1 c_2 \bar{c}_2 (\alpha_1 \delta_1 - \beta_1 \gamma_1) [G_1 + G_2]}{P_1(E_1, E_2) P_2(E_2)} \end{aligned}$$

$$G_1 := (\gamma_2 - \delta_2) c_1 (\alpha_2 \alpha_1 c_2 + \beta_2 \beta_1 \bar{c}_2)$$

$$G_2 := (\alpha_2 - \beta_2) \bar{c}_1 (\gamma_2 \gamma_1 c_2 + \delta_2 \delta_1 \bar{c}_2)$$

Using a similar proof strategy one can show that: (a) $P_1(C_2 | E_1, E_2) = P_2(C_2^* | E_2)$ if and only if $\alpha_1 \delta_1 = \beta_1 \gamma_1$ or (ii) $\alpha_2 = \gamma_2$ and $\beta_2 = \delta_2$; (b) $P_1(C_1 | E_1, E_2) = P_3(C_1^* | E_1)$ if and only if (i) $\alpha_2 \delta_2 = \beta_2 \gamma_2$ or (ii) $\alpha_1 = \beta_1$ and $\gamma_1 = \delta_1$; and (c) $P_1(C_2 | E_1, E_2) = P_3(C_2^* | E_1)$ if and only if (i) $\alpha_2 \delta_2 = \beta_2 \gamma_2$ or (ii) $\alpha_1 = \gamma_1$ and $\beta_1 = \delta_1$ (proofs omitted).

It follows then that $P_1(C_1 | E_1, E_2) = P_2(C_1^* | E_2) = P_3(C_1^* | E_1)$ if (1) $\alpha_1 \delta_1 = \beta_1 \gamma_1$ and $\alpha_2 \delta_2 = \beta_2 \gamma_2$, or (2) $\alpha_1 =$

Table 3: Frequency of participants' reposes around five focal points

	Full algebra				Sequential algebra			
	$P(C_1 E_i)$	$P(C_2 E_j)$	$P(C_1 E_i, E_j)$	$P(C_2 E_i, E_j)$	$P(C_1 E_i)$	$P(C_2 E_j)$	$P(C_1 E_i, E_j)$	$P(C_2 E_i, E_j)$
<i>AaO</i>								
'15%'			0	2			3	3
'20%'			2	6			2	0
'50%'			14	13			11	14
'60%'			4	3			2	5
'70%'			3	3			5	5
<i>SbS1</i>								
'15%'	5	3	0	1	1	2	0	1
'20%'	0	3	0	2	0	0	5	5
'50%'	15	19	14	14	23	22	13	18
'60%'	3	0	4	3	3	1	9	7
'70%'	8	4	5	4	5	6	3	6
<i>SbS2</i>								
'15%'	4	9	4	7	3	4	1	1
'20%'	10	11	2	4	9	9	0	1
'50%'	15	12	12	13	17	20	13	17
'60%'	2	3	3	3	5	4	1	0
'70%'	5	3	10	4	0	0	15	12

β_1 and $\gamma_1 = \delta_1$, or (3) $\alpha_2 = \beta_2$ and $\gamma_2 = \delta_2$. Similarly, $P_1(C_2 | E_1, E_2) = P_2(C_2^* | E_2) = P_3(C_2^* | E_1)$ if (1) $\alpha_1 \delta_1 = \beta_1 \gamma_1$ and $\alpha_2 \delta_2 = \beta_2 \gamma_2$, or (2) $\alpha_1 = \gamma_1$ and $\beta_1 = \delta_1$, or (3) $\alpha_2 = \gamma_2$ and $\beta_2 = \delta_2$. Therefore, the order is not important and one can decompose a full CBN in smaller ones while preserving the same probability distributions if (1) $\alpha_1 \delta_1 = \beta_1 \gamma_1$ and $\alpha_2 \delta_2 = \beta_2 \gamma_2$; or (2) $\alpha_1 = \beta_1$, $\gamma_1 = \delta_1$, $\alpha_2 = \gamma_2$, and $\beta_2 = \delta_2$; or (3) $\alpha_2 = \beta_2$, $\gamma_2 = \delta_2$, $\alpha_1 = \gamma_1$, and $\beta_1 = \delta_1$; or (4) $\alpha_1 = \beta_1 = \gamma_1 = \delta_1$; or (5) $\alpha_2 = \beta_2 = \gamma_2 = \delta_2$. (4) and (5) make E_1 and E_2 respectively fully undiagnostic with respect to C_1 and C_2 , which violates the faithfulness condition (see Neapolitan, 2003). (1) implies that C_1 and C_2 are conditionally independent given E_1 and that they are also conditionally independent given E_2 , that is, learning E_1 makes C_1 and C_2 independent and learning E_2 makes C_1 and C_2 independent. (2) and (3) both entail (1) and are more specific versions of (1).

Appendix B

Here we show that there are no order effects when the causes are mutually exclusive and exhaustive, i.e. when $P(C_1, C_2) = 0$ and $P(C_1) + P(C_2) = 1$. We model mutually exclusive and exhaustive causes with one node, C , that has two values: C_1 and C_2 .




Figure 4: CBN with mutually exclusive and exhaustive causes

$$\begin{aligned}
 P_4(C = C_1) &= c & , & & P_4(C = C_2) &= \bar{c} \\
 P_4(E_1 | C_1) &= \alpha_1 & , & & P_4(E_1 | C_2) &= \beta_1 \\
 P_4(E_2 | C_1) &= \alpha_2 & , & & P_4(E_2 | C_2) &= \beta_2
 \end{aligned} \tag{3}$$

Splitting the CBN from Figure 4 we get two CBNs:

$$P_5(C = C_1) = c \quad , \quad P_5(C = C_2) = \bar{c}$$




Figure 5: 'Split' CBN from E_1 to E_2

$$\begin{aligned}
 P_5(C^* = C_1^*) &= P_4(C_1 | E_1) & , & & P_5(C^* = C_2^*) &= P_4(C_2 | E_1) \\
 P_5(E_1 | C_1) &= \alpha_1 & , & & P_5(E_1 | C_2) &= \beta_1 \\
 P_5(E_2 | C_1^*) &= \alpha_2 & , & & P_5(E_2 | C_2^*) &= \beta_2
 \end{aligned} \tag{4}$$

Theorem 2. $P_4(C_1 | E_1, E_2) = P_5(C_1^* | E_2)$ when $P_{4,5}(C_1^{(*)}, C_2^{(*)}) = 0$ and $P_{4,5}(C_1^{(*)}) + P_{4,5}(C_2^{(*)}) = 1$.

Proof.

$$\begin{aligned}
 P_4(C_1 | E_1, E_2) &= \frac{P_4(C_1) P_4(E_1 | C_1) P_4(E_2 | C_1)}{\sum_C P_4(C) P_4(E_1 | C) P_4(E_2 | C)} \\
 &= \frac{c \alpha_1 \alpha_2}{c \alpha_1 \alpha_2 + \bar{c} \beta_1 \beta_2} \\
 P_5(C_1^* | E_2) &= \frac{P_5(C_1^*) P_5(E_2 | C_1^*)}{\sum_{C^*} P_5(C^*) P_5(E_2 | C^*)} \\
 &= \frac{P_4(C | E_1) P_4(E_2 | C)}{\sum_C P_4(C | E_1) P_4(E_2 | C)} \\
 &= \frac{J \alpha_2}{J \alpha_2 + (1 - J) \beta_2} \\
 J &:= \frac{c \alpha_1}{c \alpha_1 + \bar{c} \beta_1}
 \end{aligned}$$

Let $\Delta_2 := P_4(C_1 | E_1, E_2) - P_5(C_1^* | E_2)$. Then

$$\Delta_2 = \frac{c \alpha_1 \alpha_2 \beta_2 \left[1 - \frac{c \alpha_1 + \bar{c} \beta_1}{c \alpha_1 + \bar{c} \beta_1} \right]}{(c \alpha_1 \alpha_2 + \bar{c} \beta_1 \beta_2) (J \alpha_2 + (1 - J) \beta_2)} = 0$$

Since $P_4(C_2 | E_1, E_2) = 1 - P_4(C_1 | E_1, E_2)$ and $P_5(C_2^* | E_2) = 1 - P_5(C_1^* | E_2)$, then given Theorem 2 it also true that $P_4(C_2 | E_1, E_2) = P_5(C_2^* | E_2)$. Similarly we get that $P_4(C_1 | E_1, E_2) - P_6(C_1^* | E_1) = 0$ and $P_4(C_2 | E_1, E_2) - P_6(C_2^* | E_1) = 0$ (proofs omitted). ■