UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Impact of sequential reports with different source dependencies

Permalink
https://escholarship.org/uc/item/8fv4w1sq

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors
Madsen, Jens Koed
Connor Desai, Saoirse Ami
Pilditch, Toby D

Publication Date
2025

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Impact of sequential reports with different source dependencies

Jens Koed Madsen[1]
(j.madsen2@lse.ac.uk)

Saoirse Connor Desai[2]
(Saoirse.ConnorDesai@uts.edu.au)

Toby Pilditch[3]
(tdpilditch@gmail.com)

[1]Department of Psychological and Behavioural Science, London School of Economics and Political Science, Houghton Street, WC2A 2AE, London, UK.
[2]School of Psychology, University of New South Wales, Randwick, Australia
[3]Department of Experimental Psychology, University College London, 26 Bedford Way, London, WC1H 0AP, UK

Abstract

How we update our beliefs when encountering new evidence is the basis of evidential reasoning. Often, this will involve weighing up multiple pieces of evidence communicated to us by several sources (i.e., testimony). However, the testimonies of multiple sources are rarely truly independent; they may have used the same data or evidence, have the same training or background, or simply be repeating the same story as another. The nature of these dependencies among our evidence items is normatively impactful on the conclusions we should draw. Here we investigate whether participants are sensitive to such complex, yet impactful influences on their reasoning. We find a general preference for source diversity that heuristically gels with normative assertions. To our knowledge, it is the first paradigm that integrates shared background, shared evidence, and corroboration in the same design. We discuss challenges with developing and testing the intricacies of this.

Keywords: Reasoning; Dependency; Belief revision

Introduction

We rely on information from other people to form our beliefs and opinions (Hahn et al., 2012). We might check the weather app in the morning to see if it will rain or read the news on societal issues. Here, we do not have first-hand evidence but must rely on the capacity and willingness of others to provide reliable reports. This makes belief updating inherently social. As Hahn and Hornikx (2016, p. 1833) point out, reasoning is essential "for humans to learn, make decisions, and interact with others". However, determining how to update our beliefs based on others' information is not straightforward.

A key consideration is the source reliability, as individuals may vary in their ability and willingness to provide accurate information. Sources who lack expertise or are untrustworthy should influence how we process the information they report. Models such as the Elaboration-Likelihood Model (e.g., Petty & Brinol, 2011) and Bayesian frameworks (e.g., Hahn et al., 2009; Harris et al., 2015; Madsen, 2019) examine how source reliability shapes belief updating.

Beyond assessing individual sources, people must also evaluate multiple reports on the same issue, which may come from independent or dependent sources. This raises challenges beyond individual source reliability, as reports from one source may be compromised or influenced by their relationships to other sources. Such source dependencies are important to consider, perhaps increasingly so due to for example online misinformation, plagiarizing and AI.

There are multiple forms of source dependencies and people may be exposed to them simultaneously (Hahn, Corner & Harris, 2016). Dependencies may be unidirectional (a source can know about the information given by another source but not vice versa), bi-directional (two sources may discuss their views with each other before giving information), or many other combinations. The structure of dependencies between the sources may influence how their reports should normatively impact the receiver's belief revision.

Dependencies can also relate to the evidence underlying reports. Evidentiary dependencies can manifest in various ways. For example, sources may rely on the same information to draw their conclusion, they may have a similar educational background that influence their interpretation of evidence, they may come from the same political party and therefore have motivations to interpret the information in similar ways, and more. When considering reports from other people, it is normatively beneficial to have two independent assessors of the same evidence if the interpretation of it is highly technical and subject to noise (e.g., two medical professionals who independently examine the same case data or arguments, see Madsen et al., 2020).

Failure to appreciate dependencies can lead people to form erroneous belief revisions. Perceived dependencies may be societally important. For example, if a climate change denier believe that official sources are structurally dependent (and unreliable) and that internet sleuths are independent (and reliable), it should normatively influence how they integrate new reports and how they seek out information. Dependency disagreement among audiences may even cause polarisation (Young et al., 2024).

Appreciating the dependence between information sources when evaluating available reports is therefore vital for evidential reasoning tasks. While several studies have studied and tested formal models of source reliability, the present

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society* ©2025 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

study contributes to the literature by looking at the multiple forms of source dependence in combination.

Dependencies and reasoning

Bayesian Belief Networks (Pearl, 1988) can formally model dependencies. These are probabilistic graphical models that represents relationships among variables using nodes and edges. In the models, nodes denote variables (e.g., sources, hypotheses, evidence) and edges (arrows) represent the relationships between those variables (the more distant the connection between two nodes, the more independent they can be considered –for a more thorough exposition, see Schum, 1994). As an illustrative example, Madsen et al. (2020) use two dependency structures: full independence between sources (Fig. 1a) and shared reliability between sources (Fig. 1b). For full independence, all witnesses report their beliefs without any connection to other witnesses. For shared reliability, witnesses share some characteristic, which may influence how they draw their conclusions (e.g., if the witnesses have received the same educational training in interpreting data).

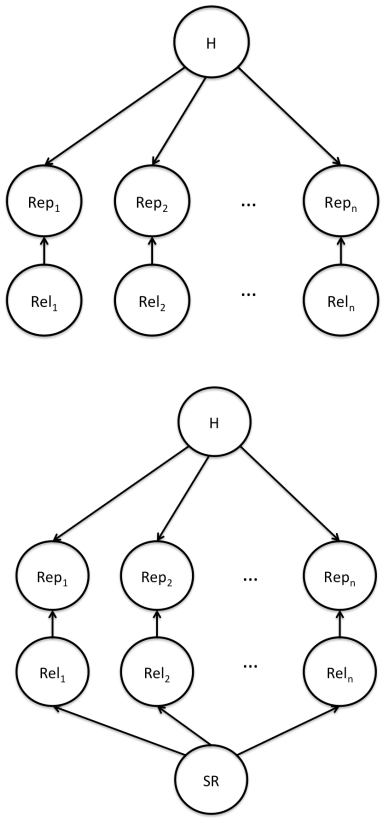

Fig. 1a: Partially independent sources of information; Fig. 1b: Sources with shared reliability

In this paper, we experimentally test how dependencies impact belief revision in a novel scenario that manipulates three forms of dependence, and their influence on subsequent belief updating about a hypothesis and the choice of potential sources. We test for the impact of sources sharing a common background/reliability (which should decrease the evidential value of sources, relative to distinct backgrounds). Second, we test the impact of sources who use the same evidence to inform their reports, which should reduce evidentiary value compared to distinct evidence bases (the degree is dictated by assumptions of noise within that evidence base, see e.g., Pilditch et al., 2019). Finally, we test the presence or absence of a unidirectional (conferring) dependence between the first and second sources (i.e., a second source may or may not have seen the first sources report) impacts judgements. To our knowledge, this is the first test of these three elements in one experimental paradigm.

Although when sources corroborate, there should be a blanket preference for (and stronger updating from) independent sources, under uncertain conditions of contradicting sources, a dependent source may be preferable (see Pilditch et al., 2020). As such, we hypothesise that participants will be aware of the benefit of reports from independent sources compared with dependent sources.

We expect this to manifest in source choices (see Methods) and in their belief revision (i.e., their posterior degree of belief should update more when reports come from independent sources). Further, we hypothesise that the use of new evidence will elicit stronger changes to posterior degrees of belief compared with reports that make use of the same evidence as the initial report (see Methods). Finally, we expect that participants will change their beliefs in the direction of the report, such that secondary reports that corroborate the first report will elicit stronger posterior degrees of belief while contradictory reports will reduce the posterior degree of belief.

In this paper we explore if people are sensitive toward source dependencies in two ways. First, if they prefer sources that are independent rather than dependent and second, if dependency descriptions impact their belief revision. As only little work tested the three forms of dependency in concert, the current paper is an exploratory study of dependency intuitions. Due to this, the experimental scenario and manipulations are kept simple. To test if participants' intuitions are qualitatively in line with these dynamics, or whether certain dependency features dominate others, we posit qualitative model predictions for their combination that follow an additive rule (i.e., more independence/diversity in factors, the stronger updates should be). We describe the experimental hypotheses in more detail in the following.

Method

We manipulate dependency of shared reliability between sources, the choice of evidence, and whether the second source corroborates or disagrees with the first source. This yields a 2 x 2 x 2 between-subjects experimental structure.

We measured three outcome variables: their belief in the hypothesis, their perception of the reliability of each source, and their preference for source choice. Having read the

instructions and consented to the experiment[1], we elicited participants' prior degrees of belief to have a baseline against which posteriors degrees of belief could be compared. To elicit the prior, the following vignette was used:

Imagine you are a fire investigator for an insurance company called to inspect a warehouse fire. You are trying to determine whether it's more likely that the fire was an ***accident*** *or an* ***act of arson****. Based on the facts of the case so far, it is safe to assume the initial probability of arson is around 50% In addition to this initial estimate, you are given two reports regarding the fire.*

Based on the above description, what do you believe the initial probability of arson is in this case?

To elicit their prior, participants responded on a scale from 0 ('it was definitely an accident') to 100 ('it was definitely an act of arson). Given the description, participants should anchor their prior around 0.5 (as seen below, participants understood the instructions, as the prior was recorded as 0.52). Thereafter, all participants read the first report to elicit the initial posterior degree of belief.

The first report comes from Eric, a forensic fire investigator who has studied causes of fires for around 30 years. Eric works at the company FireInsure. Amongst other things, Eric looks at an oil-soaked rag that was found at the warehouse. Looking at the type of oil on the rag as a possible cause for the fire, Eric reports that the fire was an act of arson.

Given Eric's report, how likely do you think now that the fire was accidental or an act of arson?

Participants responded on the same 0-100 scale as for the prior belief, allowing us to test the belief revision given the first report on the potential cause of the fire. Finally, we tested the effect of experimental conditions on the perceived reliability of sources of information by asking participants to evaluate Eric's reliability:

How reliable do you think Eric the forensic fire investigator is?

Participants rated source reliability on a scale from 0 ('Eric is completely unreliable as a source') to 100 ('Eric is completely reliable as a source')[2].

All participants read the initial description and the first report from Eric. After this, we employ a 2x2x2 between-subjects design (manipulating shared background, shared evidence, and whether the source corroborates with the original report). Participants all saw another set of reports from two different sources: Robert and Jimmy. These expert sources also report on the possible cause of the fire. These sources all claim to know something about the cause of the fire. For all conditions, Robert saw the first report from Eric (the original source) before giving their own report while Jimmy did not to see Eric's report before giving their report. This manipulates the dependency between Eric and the two available sources where Robert's report is dependent on Eric's while Jimmy's is independent.

We compared if sources come from the same or different companies (shared background) as Eric, if they make use of the same of different evidence (shared evidence) as Eric, and if they corroborate or disagree with Eric's assessment (corroboration). Before seeing Robert and Jimmy's reports, participants were asked which of the sources they would query if they could. The following is an example of the condition where Robert and Jimmy come from *different* companies and make use of *different* evidence to arrive at their report.

In addition to Eric, there are two other people who can give reports on the cause of the fire.

Robert is also a forensic fire investigator who has studied causes of fires for around 30 years. Robert works at a different company, BurnOut and looks at the electrical wiring as a possible cause for the fire. Robert has been able to see the Eric's report before making his own report.

Jimmy is also a forensic fire investigator who has studied causes of fires for around 30 years. Jimmy works at a different company, FlameSafe and looks at the gas pipes as a possible cause for the fire. Jimmy has NOT been able to see Erics report before making his own report.

If you could choose to see either Robert's or Jimmy's report on the cause of the fire, which one would you choose?

Participants were asked if they would choose either Robert, or Jimmy's reports, or if they believe reports from the two sources would be of equal value. We later refer to this measure as the reported *source choice*. This measure allows us to explore whether participants were sensitive to the dependency (whether the source has seen Eric's report). After this choice, participants see Robert's report. The following is an example of a corroborating report.

Imagine you read Robert's report (who works as the same company, looks at the electrical wiring as a cause of the fire, and has seen Eric's report before making his own report).

[1] To ensure that participants read the instructions, two comprehension checks were carried out. One asked them to relay the task of belief revision and the other asked them to relay their task on reliability perception. If participants failed either, the experiment ended (they still received full payment for participation despite failing the comprehension checks to thank them for their time).

[2] To have similar interpretations of 'reliability', we added the following note to the instructions: 'for the purpose of the survey, reliability can be understood as the capacity and willingness to provide relevant information about a given topic.'

Robert's report corroborates Eric's, and he also reports that the fire was an act of arson.

Given Robert's report, how likely do you now think that the fire was accidental or an act of arson?

Note: in this case, remember you have only seen Eric and Robert's reports.

Participants responded on the same 0-100 scale as for the prior belief, allowing us to test the belief revision given the first report on the potential cause of the fire. In addition, they were asked to consider Eric and Robert's reliability after seeing Robert's report. Such a measure allows us to test how people update their beliefs given the new report. Finally, participants are given Jimmy's report. Here, they are asked specifically to consider how they would respond if they had only seen Eric and Jimmy's report (and not Robert's). The following is an example of a corroborating report.

Imagine instead that you read Jimmy's report (who works as the same company, looks at the gas pipes as a cause of the fire, and has seen Eric's report before making his own report). Jimmy's report corroborates Eric's, and he also reports that the fire was an act of arson.

Given Jimmy's report, how likely do you now think that the fire was accidental or an act of arson?

Note: in this case, remember you have only seen Eric and Jimmy's reports.

As with after Robert's report, participants were asked to consider the potential cause of the fire (from 0-100) as well as the reliability of Eric and Jimmy.

Participants

We recruited 480 participants from Prolific. To participate in the experiment, participants had to be Native English speakers with an age of 18+. In total, 35 participants were discarded for failing to respond to the comprehension check, leaving 465 participants for the analysis. Mean participant age was 43.85 (SD = 13.66). 230 identify as male, 232 as female, 1 a non-binary, and 2 did not want to say. Median completion time was 4.7 minutes (a priori estimated time of completion was 5 minutes) and participants were paid £0.75, yielding an effective pay of roughly £9.57 per hour.

Results

All analyses were conducted using JASP (Version 0.18.3). We report Bayesian analyses of the data which are interpreted according to the convention proposed by Lee and Wagenmakers (2014). This convention provides a graded scale for interpreting the strength of evidence in favor of the alternative hypothesis, ranging from 'anecdotal' to 'extreme'. Specifically, a Bayes Factor (BF^{10}) of > 100 is considered to provide 'extreme' evidence for the alternative hypothesis.

We examined whether people recognized dependencies and how this influenced belief revision. We assessed sensitivity through source choices, determining if participants recognized Robert's dependence on Eric, having seen his report, while Jimmy remained independent, forming his conclusion without prior exposure.

Source choices. Fig. 2 shows a strong overall preference for an independent source. The preference is stronger when shared reliability, and shared evidence are the same. However, when evidence features were different, dependency was less important, and more participants say they would prefer to hear from dependent sources or think that independent and dependent reports are the same.

Bayesian contingency table analyses was performed to examine the relationship between source choice and evidence features (shared reliability, shared evidence, corroboration). The results of the Bayesian analysis for shared reliability, shared evidence, and corroboration provided extreme

Fig. 2: Source choice as a function of evidence features.

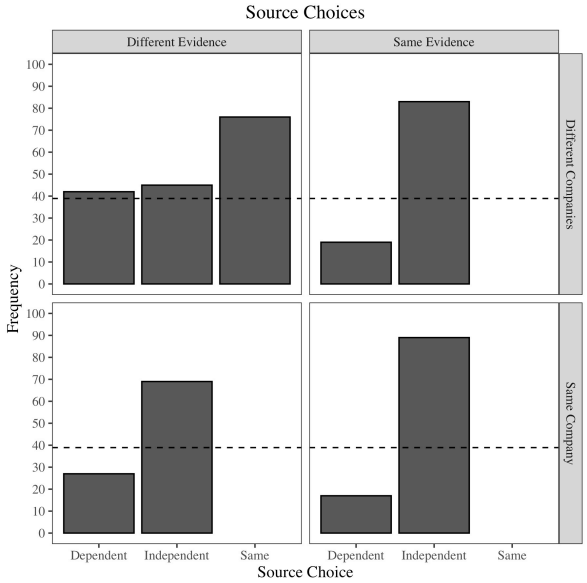

Note: Figure illustrates the frequency of choices when participants were asked whether they would consult dependent (Robert), independent (Jimmy) sources, and the number of participants who said they were equivalent.

evidence for the alternative hypothesis over the null hypothesis, with a BF_{10} > 100,000, which is extremely strong evidence for the alternative hypothesis as compared with the null hypothesis. As such, the data indicates that people are sensitive to the benefit of epistemic independence. That is, they broadly prefer sources that are unconnected to the original source rather than a supposed benefit of another person looking over the same evidence.

The only condition where this pattern does not emerge is in the condition where Robert and Jimmy are described as from different companies than Eric and where they make use of different evidence to reach their conclusion. In this case, participants are more likely to believe the reports are of equal

value, suggesting participants do not perceive structural dependence to be as impactful when there is sufficient diversity from other sources (reliability, evidence).

Prior and Posterior belief in hypothesis. In addition to the source choices, we examined whether awareness of the importance of dependency impacts belief revision upon seeing the reports. Figure 3 shows the prior belief in the hypothesis, and the posterior belief in the hypothesis after seeing the first source's report, and then after seeing an additional report from either a dependent or independent source. A Bayesian Repeated Measures ANOVA was conducted to examine the effects of source (dependent or independent), shared reliability, shared evidence, and corroboration on belief in the hypothesis at different timepoints. That is, prior to observing any evidence, and then after seeing reports from the various sources. The model included four levels of the within-subjects factor timepoint. Looking at the analysis of effects, there was extreme evidence for an interaction between the second source (independent and dependent) and corroboration, an effect of the second source, and an effect of corroboration, all with a BF^{10}s > 100,000. For the remaining effects, there was evidence in favour of the null, BF^{10}s < 1.

Belief revision results were mixed. The corroboration or disagreement manipulation yielded differences (they reduced their degree of belief in arson as the cause of the fire when

Fig. 3 Belief in hypothesis before and after reports as function evidence features.

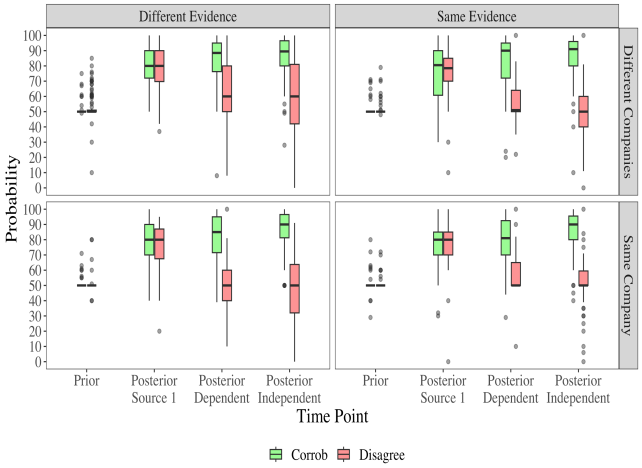

Note: Boxplot showing median belief in the hypothesis before reading reports (prior), posterior belief in the hypothesis after seeing the first source's report (Eric), then posterior after seeing the first source and the dependent source (Robert), and the posterior after seeing the first source and the independent source (Jim).

the second reports disagreed with Eric's conclusion). However, the experimental manipulation of dependency was not sufficient to change their posterior degree of belief in the hypothesis when they compared reports from Robert and Jimmy. This suggests that, while participants preferred to see reports from an independent source in their source choice, they did not update differently when seeing the independent report. This may be due to the experimental design where all participants saw both Robert and Jimmy's report and were asked to imagine that they had not seen Robert's report.

Source reliability. A Bayesian Repeated Measures ANOVA was conducted to examine the effects of shared reliability, shared evidence, and corroboration on the reliability of the dependent and independent sources. The analysis of effects showed extreme evidence for the interaction between second source and corroboration, BF_{10} > 100,000, such that corroboration generally leads to increases in perceived reliability, whilst contradiction leads to a penalty. There was also moderate evidence for the effect of source (independent vs dependent), BF_{10} = 9.937, and extreme evidence for the effect of shared evidence, BF_{10} = 236.8. The remaining effects showed evidence for the null, BF_{10}s < 1. Fig 4. confirms the analysis and shows that independent sources generally suffer greater penalties to perceived reliability because of contradiction than dependent sources. This goes against our predictions, as independent sources should be seen as *more* reliable than dependent

Fig 4. Reliability of dependent and independent sources as a function of evidence features.

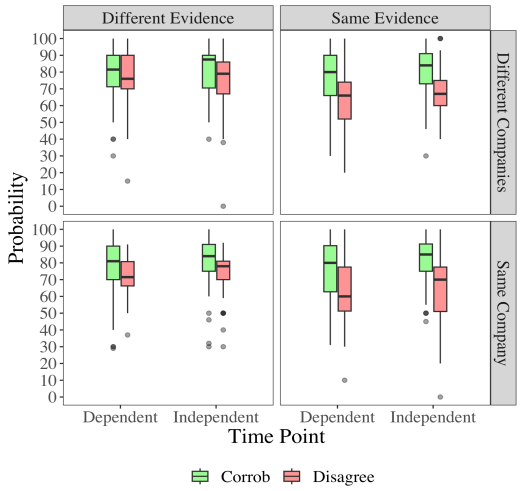

Note: Boxplot showing median reliability judgments for dependent and independent sources. The biggest difference in reliability is seen when independent and dependent sources are from the same company and give the same evidence.

sources.

Discussion

In this paper, we explored if people are sensitive toward source dependencies. Specifically, we test a novel scenario that manipulates whether sources share the same background, the same evidence, and whether they corroborate or disagree with the original statement. We found a strong preference for

source diversity, with importance decreasing as we move away from the core report content. That is, corroboration compared with contradiction was the most important factor, followed by the distinction between dependent and independent secondary sources. Shared evidence held less weight, while shared reliability was of least concern. Our findings suggest a clear hierarchy in source diversity preferences. Importantly, the results suggest that people know that relationships between sources and the evidence they use is important when they consider who and what to believe. This indicates that people prefer independent sources when asked but may – in the scenario we used – lack sensitivity to identity these dependencies. That is, this does not appear to translate directly into belief revision regarding the hypothesis or source reliability.

The relationship between sources and the evidence they use to form their reports is central to reasoning. If a person receives two reports about the same issue, it is essential to know whether the two sources have based their conclusions on the same evidence or if they have reached their inferences independently. If a person believes reports are independent where they are strongly dependent, it can lead to reasoning errors and poor decision-making. While formal models of dependency have been around for a while (e.g. Bovens & Hartmann, 2003), it is only recently that structured empirical testing of people's capacity to understand this problem has emerged (e.g. Madsen et al., 2020). This paper adds to the growing amount of literature on this topic.

Although modest, our results suggest a disconnect between people's preferences and their ability to identify dependencies between sources and incorporate these dependencies into their reasoning. Our findings are consistent with other studies which have shown that although people can in principle identify dependencies this does not translate into measures of belief updating (e.g., Yousif et al., 2019). Our findings highlight the importance of further investigation into the importance of source dependencies on belief revision.

On specific limitation was the topic participants considered – the cause of a fire, which is presumably an unusual event for most participants. As dependencies may be technical and slightly convoluted, it is possible that they have greater impact on belief revision when they concern issues that are more immediate to the participants. For example, Cosmides and Tooby (1992) showed that people performed better on a naturalistic version of the Wason Selection Task. People may be better able to translate their preference for independent information to their belief revision if they can identify independence in a more familiar setting, such as when considering gossip or workplace rumours (e.g., Mercier & Miton, 2019). A crucial area for future research is to investigate whether these results generalize to other scenarios.

Concluding remarks

We show that participants have a general preference for diversity among sources. All things being equal, people seem to prefer reports from sources that are independent. Yet, this awareness of dependency weakens the further away from the report content you get. In the study, results suggest that whether the new sources corroborate or contradict the initial report is important (closest), then whether the sources have been able to see the initial report is next most (also next closest structurally), then whether the sources share evidence when reaching their conclusions (weaker still, and further still), and whether they have shared reliability such as coming from the same company (remotest of all).

Our findings suggest people are aware of the importance of dependencies. This has both theoretical and practical implications. On the theoretical level, the study provides evidence that people's intuitions are in line with normative assumptions of the importance of source independence. This indicates that people engage with the sources in a reasonable way (Madsen et al., 2024). On the practical side, the role of perceived dependencies may be important for many societal debates. It may frame how people search for and treat information about climate change. Climate change deniers may believe that official sources are structurally dependent (and unreliable) and that internet sleuths are independent (and reliable). If so, they should seek out information from perceived independent sources while reports from official sources should have less impact on their belief revision. This implication of dependencies is in line with studies that show how Bayesian networks can yield climate change polarization (Cook & Lewandowsky, 2016). As such, source dependencies offer a complementary perspective to the role of perceived reliability when studying why some people are hesitant to update their beliefs by seeking new evidence.

References

Bovens, L., & Hartmann, S. (2003). *Bayesian epistemology*. Oxford University Press

Cook, J., & Lewandowsky, S. (2016). Rational irrationality: Modeling climate change belief polarization using Bayesian networks. Topics in Cognitive Science, 8(1), 160–179.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. *The adapted mind: Evolutionary psychology and the generation of culture*, *163*, 163-228

Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*. *11*(2), 77-83

Hahn, U., Harris, A. J., & Corner, A. (2016). Public reception of climate science: Coherence, reliability, and independence. *Topics in cognitive science*, *8*(1), 180-195.

Hahn, U., & Hornikx, J. (2016). A normative framework for argument quality: Argumentation schemes with a Bayesian foundation. *Synthese*, *193*(6), 1833-1873

Hahn U., Harris A. & Corner A. (2009). Argument content and argument source: An exploration, *Informal Logic* 29, 337-367.

Hahn, U., Oaksford, M., Harris, A. J. L. (2012). Testimony and argument: A Bayesian Perspective. In F. Zenker (Ed.), *Bayesian argumentation: The practical side of probability* (pp. 15–38). Springer Science + Business Media

Harris, A. J. L., Hahn, U., Madsen, J. K., & Hsu, A. S. (2015). The appeal to expert opinion: Quantitative support for a Bayesian network approach. *Cognitive Science*, *40*(6), 1-38

Madsen, J. K. (2019). Voter reasoning bias when evaluating statements from female and male political candidates. *Politics & Gender*, *15*(2), 310-335.

Madsen, J. K., De-Wit, L., Ayton, P., Brick, C., De Moliere, L. & Groom, C. J. (2024) Behavioural science should start by assuming people are reasonable, *Trends in Cognitive Science*

Madsen, J. K., Hahn, U., & Pilditch, T. D. (2020). The impact of partial source dependence on belief and reliability revision. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 46* (9), 1795–1805.

Madsen, J. K. & Wong, M. (2023) Comparing predictions from the Elaboration Likelihood Model and a Bayesian model of argumentation, *Proceedings of the 45th Annual Conference of the Cognitive Science Society*

Petty, R. E., & Briñol, P. (2011). The elaboration likelihood model. *Handbook of theories of social psychology*, *1*, 224-245.

Pilditch, T. D., Hahn, U., Fenton, N., & Lagnado, D. (2020). Dependencies in evidential reports: The case for informational advantages. *Cognition*, *204*, 104343.

Pilditch, T.D., Hahn, U., & Lagnado, D. (2019). Shared Evidence: It all depends... In Goel, A., Seifert, C., & Freska, C. (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 2571-2577). Austin, TX: Cognitive Science Society.

Schum, D. A. (1994) *The Evidential Foundations of Probabilistic Reasoning.* New York: John Wiley & Sons

Young, D., de-Wit, L. H. & Madsen, J. K. (2024) Bayesian Belief Polarization due to Differential Perceptions of Source Independence, *Proceedings of the 46th Annual Conference of the Cognitive Science Society*