Reasoning about dissent: Expert disagreement and shared backgrounds

Jens Koed Madsen (jens.madsen@ouce.ox.ac.uk)
School of Geography and the Environment, University of Oxford
OX1 3QY, South Parks Road, Oxford, United Kingdom
Orcid: 0000-0003-2405-8496

Ulrike Hahn (u.hahn@bbk.ac.uk)
Department of Psychological Sciences, Birkbeck College, University of London
WC1E 7HX, Malet Street, London, United Kingdom

Toby Pilditch (t.pilditch@ucl.ac.uk)
School of Geography and the Environment, University of Oxford
OX1 3QY, South Parks Road, Oxford, United Kingdom
&
Department of Experimental Psychology, University College London
26 Bedford Way, WC1H 0AP, London, United Kingdom

Abstract

Sequential testimonies where more or less reliable sources argue about an issue are central to public debates. Often, the majority of sources may argue that a hypothesis is true while a minority dissenter may claim the opposite (e.g. scientists and lobbyists in the climate change debate).

In this paper, we show that people are sensitive to source reliability as well as the structural relationship between the sources. Participants follow Bayesian predictions for revising belief in the hypothesis *and* the reliability of the competing sources given majority consent, minority dissent, and shared reliability between sources. Shared reliability and dissent is a key issue for public debate and belief revision. The paper provides novel insight into the workings of these aspects.

Keywords: Source reliability; Shared reliability; Source dependency; Bayesian modelling; Belief revision

Introduction

Information is crucial to revising or maintaining beliefs, to making decisions in an uncertain world, and to compare and contrast support for competing hypotheses. While we can certainly acquire information through personal experience (e.g. witnessing congested traffic may change the route we travel to work, participating in a public demonstration may give an impression of the degree of support for a particular cause, etc.), most of the information we get in our everyday lives comes via other people. Meteorologists provide us with necessary weather information for planning the day, news readers give us an overview of relevant events that happen within our respective countries and abroad, and friends, family members, and co-workers provide invaluable information on a range of issues that help us appreciate their lives, consider information we have not been privy to before talking to that person, or information that is necessary for doing our respective jobs.

Appeals to authority have traditionally been regarded as a reasoning fallacy – this is due to the fact that perceived authority should not add credence to the conditional link between the evidence and a hypothesis. That is, whether or not a piece of information increases the likelihood of a hypothesis is, in principle, independent from the source that conveys the piece of information. Classically, this has led some people to be sceptical of appeals to authority.

The notion that appeals to authority should be distrusted in principle reverberates in theories of argumentation and reasoning. For example, two prominent models of persuasion, the Elaboration Likelihood-Model (Petty & Cacioppo, 1984) and the Heuristic-Systematic Model (Chaiken & Maheswaran, 1994), classify appeals to authority as a shallow and weak cue. In this view, people should disregard the characteristics of the source as they are given greater incentive to interrogate and elaborate on the evidence and its relation to the hypothesis. In other words, as the incentive to understand the link between evidence and claim increases, the nature of the source should matter less and less.

While it is it true in principle that the messenger neither adds nor subtracts to the link between evidence and claim, overlooking the epistemic impact of perceived source reliability neglects a crucial communicative and reasoning function. In a world where sources can lie and make up evidence, their reliability becomes crucially linked with the strength of the argument. Additionally, in a highly uncertain world, some information requires deep expertise to process (e.g. climate data may be accessible to a general population, but requires considerable expertise to adequately model and understand). Given the capacity to misinform and generate mistaken causal models due to a lack of expertise, the reliability of the speaker is an important element for people to update and revise their beliefs in the world.

In line with this perspective, the impact of the perceived reliability of a source is shown to be crucial for reasoning and decision-making. Treating the reliability of a source as a shallow cue, the literature on persuasion has shown the impact of appeals to authority (Petty & Cacioppo, 1984; Tormala & Clarkson, 2007), the developmental literature suggests children seek out credible figures to guide their perception of the world (Harris & Corriveau, 2011), and appeals to authority have been shown to impact legal reasoning (Lagnado et al., 2013). Further, it increases adherence with persuasion strategies (Cialdini, 2007), and perceived reliability is able to predict whether or not people believe an unknown policy is good, given recommendations from different political sources (Madsen, 2016).

The paper explores three aspects of perceived reliability. First, it replicates a Bayesian model of the impact of sequential reports from more or less reliable sources. This replication shows people update their beliefs in a hypothesis given reports from sources *as well as* updating the perceived reliability of the sources themselves. Second, it replicates recent findings that shared reliability (e.g. sources sharing a common background) impacts the degree of belief in a hypothesis and the perceived reliability of sources, in line with Bayesian predictions. This explores aspects of source dependency. Finally, we extend this work by presenting novel findings on the impact of minority dissenters on belief in a hypothesis and the perception of reliability among sources, given the introduction of shared reliability. Minority dissent and shared reliability are crucial aspects of information transmission (see see Whalen et al., 2013 for dependency and Perfors et al., 2018 for minority dissent), as they appear a number of domains – indeed, most debates are characterised by sources that disagree. For example, in climate change both are apparent and important factors of public debate.

A Bayesian approach to source reliability

Whilst some have argued reliance on the reliability of others to revise subjective beliefs about the world is a shallow persuasive cue (Petty & Cacioppo, 1984; Chaiken & Maheswaran, 1994), others have argued reliance is rationally justified and a necessary component of belief revision (see Bovens & Hartmann, 2003; Hahn et al., 2009).

The latter applies a Bayesian perspective to reliability. Bayesian reasoning uses subjective, probabilistic degrees of belief in propositions where Bayes' theorem integrates prior beliefs with the likelihood ratio to estimate the posterior degree of belief (Howson & Urbach, 1996). Bayes is an alternative to logicist approaches to reasoning (Oaksford & Chater, 1991) and has been applied to argumentation theory (Hahn & Oaksford, 2006; 2007), which has found Bayesian reasoning can account human information integration in practical reasoning (see Oaksford & Chater, 2007).

The Bayesian approach suggests that people's subjective perceptions of the reliability of the speaker normatively should yield different information integration. For example, if the messenger has no expertise, the information may be regarded as pure noise (as it is equally likely to be true or false). In this case, the recipient should not revise her beliefs one way or another. Comparatively, if low trustworthiness entails simple misinformation, the recipient may increase her belief in the opposite direction given positive reports from a distrusted source. Due to the Bayesian nature of the above models, the reliability function of reports relies on conditional probabilities (see e.g. Madsen, 2016 where participants revise their beliefs negatively in a proposed policy given positive reports from subjectively distrusted politicians).

More formally, the model integrates two components to account for overall reliability: perceived trustworthiness and perceived expertise (Hahn et al., 2009)[1]. In this framework, expertise refers to the *capacity* to provide accurate information about the topic in question. This is highly domain-dependent. For example, an astrophysicist may be able to calculate the mass of a distant celestial body, but may not be able to give a valid economic forecast. While expertise refers to capacity, trust refers to the *intention* of providing true and accurate information to the best of ones ability. For example, the astrophysicist may omit data points that contradict personally held theories or beliefs. The model components are orthogonal, as a person can be highly expert in some domain while at the same time be entirely untrustworthy – or vice versa. The orthogonal assumption is theoretically grounded (Bovens & Hartmann, 2003) and empirically supported (Harris et al., 2015)

Formally, Bayes' theorem is used to integrate reliability where the posterior degree of belief in the hypothesis (H) given the representation (Rep) yields:

$$\mathrm{P(H|Rep)} = \frac{P(H)x\,P(Rep|H)}{P(H)x\,P(Rep|H) + P(\neg H)x\,P(Rep|\neg H)}\ ^{2}$$

The formalisation predicts how people should integrate uncertain information from more or less reliable sources. Model predictions have enjoyed a good fit with behavioural data (Harris et al., 2015; Madsen, 2016). Overall, the findings suggest people are sensitive to the reliability of the individual speaker and integrate the information from the speaker in a normatively rational manner.

Shared reliability: corroboration and negation

The empirical work underpinning the Bayesian source reliability model suggests that people do modulate information integration given perceived speaker reliability. The influence of reliability on belief revision means that the perceived source reliability *itself* is important in the belief revision process.

If the recipient believes the source is credible, she should revise her beliefs more positively if the source provides positive reports for a hypothesis. As a consequence, if the perceived source reliability changes, Bayesian (normative) models entail that the effect of this source should change for future reports and the impact of the already observed report. That is, if a speaker is revealed to be less than credible, audiences should be more likely to disregard any reports from that source in the future. Changes to the perceived

[1] The operationalization of reliability as an amalgamation of perceived expertise and trustworthiness is remarkably close to findings in social psychology where reliability is defined as an amalgamation of traits related to *warmth* and *competence* (Fiske et al., 2007; Cuddy et al., 2011).

[2] P(Rep|H) = P(Rep|H, Exp, T) * P(Exp) * P(T) + P(Rep|H, ¬Exp, T) * P(¬Exp) * P(T) + P(Rep|H, ¬Exp, ¬T) * P(¬Exp) * P(¬T) + P(Rep|H, Exp, ¬T) * P(Exp) * P(¬T); mutatis mutandis for P(Rep|¬H)

reliability of sources can conceivably happen for a number of reasons – for example if the source corroborates a highly unlikely hypothesis.

Further, reports are seldom made in isolation. Frequently, people will see multiple reports for a given issue. These sources, perceived as more or less credible by the recipient in question, may argue for or against a hypothesis. For example, in climate change debates, pundits, experts, and members of the media frequently make predictions about a particular hypothesis or issue. Considering flood risks in coastal areas, many experts tend to warn that weather will become more extreme and floods more prominent (e.g. in Miami). However, minority dissenters may argue that floods will not change over time. Here, we have multiple sources (e.g. scientists) that corroborate and support a hypothesis (rising floods) and a dissenter (e.g. a senator) who negates the hypothesis.

People's prior belief in the hypothesis, their perceived reliability for each source, and their perception of the dependency of sources (e.g. shared reliability versus independent sources) should normatively influence their belief in the hypothesis and perceived reliability, given positive or negative reports from the sources. The paper explores whether this is the case empirically.

In order to approach these questions, we use Bovens and Hartmann (2003) foundational and Bayesian perspective on modeling source reliability. Aside from suggesting people should revise their belief in the reliability of the source and in the hypothesis given sequential reports, their models show that *the structure* of the perceived relationship of sources influence the degree to which their reports should these beliefs given multiple testimonies.

Figure 1a-b illustrates different structural relationships between independent sources with independently perceived reliability ($Rel_{1\text{-}4}$) who provide a report ($Rep_{1\text{-}4}$) concerning a hypothesis (H). 'Independent sources' refer to situations where the sources can be considered entirely independent of one another (Fig. 1a). For example, climate scientists may run studies independently of each other and report their findings with no knowledge of the findings of other scientists (here, the strength of the report will in part depend on each reports personal reliability).

Comparatively, if sources share a common background (e.g. the scientists may have been trained at the same school to use a specific model to explore climate phenomena), they become partially dependent (Fig. 1b). In this case, the sources have shared reliability (SR).

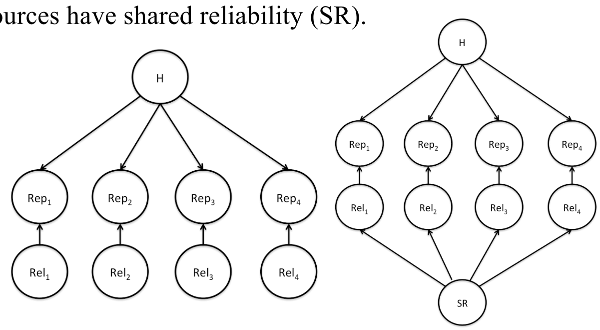

(1a) *(1b)*

Fig. 1a-b: independent sources and sources with shared reliability

Shared reliability constrains the informativeness of a source, as their reliabilities are influenced by the common-cause (e.g. attending a good or bad school). That is, the common background can weaken the impact of the reports provided by these sources. More intuitively, in finding out that sources share a compromising background (e.g. have all attended a fraudulent school), then the individual reliabilities of those sources are compromised, and in turn the strength of their support. Bovens and Hartmann (2003) provide a formal way to calculate "…how the posterior probability of the reliability of the nt^h witness increases as more and more witness reports from in" (p. 79):

$$P^{*(n)}(REL_n) = P(REL_n|REP_1, \ldots, REP_n)$$
$$= \frac{h[us(s+a\bar{s})^{n-1} + \bar{u}t(t+a\bar{t})^{n-1}]}{h[u(s+a\bar{s})^{n} + \bar{u}(t+a\bar{t})^{n}] + \bar{h}[u(s+a\bar{s})^{n} + \bar{u}(t+a\bar{t})^{n}]}$$

where u is the probability of the shared background being reliable, P(SR) – that is, how reliable the source is seen to be prior to any information about shared reliability, s is the conditional probability: $1 > P(Rel_i|SR)$ – that is, the likelihood that source i is reliable *given* the shared reliability. The conditional probability $> P(Rel_i|SR) > 0$ is represented by t, whilst a is a randomization parameter (that is, the degree of noise), and h is the prior probability of the hypothesis (that is, degree of belief in the hypothesis prior to any reports).[3]

In sum, the equation shows that the posterior degree of reliability of the n^{th} witness (or source) depends on the randomization parameter (a) and prior probability of the hypothesis (h). For example, if $a = .9$ and $h = .3$, initial witness reliability falls from .5 to .25 (see p. 80), but increases as additional positive reports confirm the initial report.[4]

Recently, Madsen et al. (2018) tested this intuition. That is, whether people update their beliefs in the reliability of the source and the belief in the hypotheses when they experience sequential corroborative testimonies. In their study, all reports *corroborated* the hypotheses (that is, all sources provided positive reports for the hypothesis). However, as mentioned, many (if not most) debates are between sources that disagree about a particular issue. For this reason, it is imperative to understand the function of (minority) dissenters.

Madsen et al. find support for the model of corroborating sources, as P(Rel) decreased given a corroborative report of an unlikely hypothesis, but subsequently increased as more corroborative reports were given. Further, when participants learned sources attended the same school, they adjusted their posterior degree of belief negatively for the hypothesis *and* source reliability. The effect was stronger if experts'

[3] C.3 (pp. 136-137) and C.4 (pp. 137-138) in Bovens and Hartmann (2003) provide the full derivation for $P^{*(n-1)}(REL_n) = P(REL_n|REL_1, \ldots, REL_{n-1})$ and $P^{*(n)}(REL_n) = P(REL_n|REP_1, \ldots, REP_n)$ respectively

[4] The current study does not elicit a randomization parameter

school was bad compared with sharing a school described as 'excellent'. Finally, their study suggests people revise posterior degree of belief in the reliability of sources retrospectively. That is, as $sources_{2\text{-}3}$ provided reports, the reliability of $source_1$ was adjusted to be in line with perceived reliability of the n^{th} source.

This paper extends this work by exploring three facets of source dependency and reliability. First, a source may corroborate or negate a report for a given hypothesis. We test how participants update their beliefs in the hypothesis and the reliability of each source given corroborative reports from $sources_{1\text{-}2}$ and a negative report from $source_3$ (denoted by '+' and '-' respectively in the Fig. 2a and b).

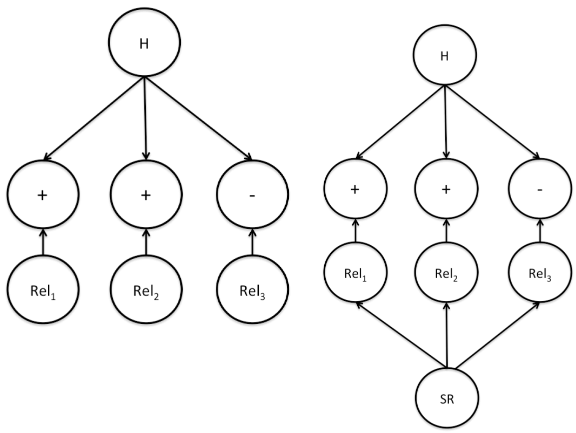

(2a) *(2b)*

Fig. 2a-b: Negation from independent sources and sources with shared reliability

Second, we explore how dependency impacts perception of the hypothesis and reliability when the 3^{rd} source dissents and negates the reports of $sources_{1\text{-}2}$. Here, we use the same a shared reliability structure (Fig. 2b).

We use the experimental design of Madsen et al. (2018), altering it to explore the following:

Given reports from the 3^{rd} dissenting source, we expect decreases in reliability for all sources and a decrease in belief in the hypothesis. This is due to the fact that dissent adds additional uncertainty to the hypothesis (the initial sources may be wrong) and onto the sources themselves (either the 3^{rd} source or the initial sources may have been mistaken/providing bad information). For the likely scenario we expect a significant drop in reliability for the 3^{rd} source in particular, as this source goes against a very likely hypothesis *and* two corroborating reports. In addition, in accordance with Madsen et al (2018), P(Rel) should *decrease* when source 1 reports an unlikely hypothesis, but subsequently recover, as source 2, still perceived to be independent, corroborates the unlikely prediction.

Method

Material and procedure: To replicate Madsen et al. (2018) and enable direct comparisons, we use their method and materials. To test model predictions, we use low and high probability scenarios. In the low probability scenario, participants were asked to evaluate the likelihood of a crash in the stock market with the following description:

> *"Imagine you are watching a news programme about the economy. Specifically, the programme considers whether or not the UK stock market will crash (i.e. fall by more than 30%) within the next 6 months. Historically, the likelihood of a crash occurring within a 6-month window is 5%.*
>
> *In your opinion, how likely is the UK stock market to crash within the next 6 months?"*

Having read this, participants provided prior estimates for their beliefs in the hypothesis on a scale from 0-1 (0: I am completely certain the stock market will NOT crash within the next 6 months; 1: I am completely certain the stock market will crash within the next 6 months). To elicit the reliability of sources, we defined reliability:

> *"Reliability can be defined as having access to relevant information about a topic, and a willingness to say what you believe to be the true state of the world.*
>
> *How reliable are economists in predicting the market crashes?"*

Having read this, participants provided their belief in the source reliability from 0-1 (0: economists are completely unreliable; 1: economists are completely reliable). Reports from sources were provided as interviews with experts on the subject. For example:

> *"Now, imagine that an economist, Robert, is being interviewed about the economy. Robert states the following: "I am completely certain the stock market will crash within the next 6 months."*
>
> Given *Robert's report, how likely is the UK stock market to crash within the next 6 months?"*

Participants then gave subjective estimates of their beliefs in the hypothesis and in each source hitherto presented. Sources were presented sequentially. To test the effect of negation, $sources_{1\text{-}2}$ always corroborated the hypothesis and $source_3$ always negated the hypothesis. This implemented a minority dissenter. The dissent only functions in light of the initial corroborations. As such, the dissenter had to be at the end of the scenario. Further, placing the dissenter towards the end allowed for replication of corroborative $reports_{1\text{-}2}$, as participants had not yet been exposed to dissent.

Finally, having seen the three sequential reports, the participants were told the sources were partially dependent (i.e. shared a background), which was manipulated between-subjects as either high or low quality (SR Condition). An example of the high quality SR Condition statement for the low likelihood scenario:

> *"It turns out, all the interviewed economists studied at the same school and subscribe to the same economic theories. Their school has a very good reputation for excellent teaching and accurate approaches to economy.*
>
> *Given the fact that they all studied at the same school and follow the same economic theories, how likely is the UK stock market to crash within the next 6 months?"*

After each report and the SR condition, P(H) and $P(Rel_{1\text{-}n})$ were measured. Participants read both scenarios in a counterbalanced order, with the SR Condition manipulated independently for each scenario.[5].

Participants: 100 participants (71 female, μ_{age} = 34.51, σ = 11.49) were recruited from the online recruitment source Prolific Academic. All had to be aged 18+ and native English speakers from either the UK or the USA. All participants had to have a prior completion rate of 95%. Median completion time was 5.56 min (σ = 2.11) and participants were paid £0.8 (resulting in an effective fair hourly wage of £8.63/hour for participation).

Results

All inferential statistics reported below were Bayesian[6], and were conducted used the JASP statistical software (JASP Team, 2018). The probability manipulations were successful in generating high and low estimates for the two scenarios: The market crash scenario was rated as unlikely (μ = .337, σ = .243) and the salmon growth scenario was rated as likely (μ = .806, σ = .116). In both scenarios, sources were rated higher ($P(Rel_{Economist})$: μ = .638, σ = .156; $P(Rel_{Biologist})$: μ = .731, σ = .128). Importantly, though, both sources were rated positively, which allows for the testing of whether positive reports of unlikely hypotheses influence reliability estimates negatively.

Following predictions from Bovens and Hartmann (2003), we expect positive reports of an unlikely hypothesis to lead to an initial decrease in estimates of reliability. To test this, we use repeated measures ANOVA ($P(Rel) – P(Rel_1|Rep_1)$). We observe a negative revision of reliability of source 1 given a positive report of an unlikely hypothesis (*N* = 100), BF_{10} = 179636.1 (in the current design, the source predicts the stock market will crash within a 6-month period). However, as participants learn another source also provides a positive report ($P(Rel_1|Rep_1) – P(Rel_1|Rep_2)$), they revise their belief in the initial source and revise reliability in a positive direction (*N* = 100), BF_{10} = 798759.1.

When a third source then contradicts ($P(Rel_1|Rep_2) – P(Rel_1|Rep_3)$), the reliability of the original reporter is then reduced once again (*N* = 100), BF_{10} = 352673.82 We further note strong evidence for a null difference in the estimated reliabilities across the three sources (*N* = 100), BF_{10} = 0.127, despite the presence of a contradicting minority, suggesting that all sources are penalized given the dissent among them.

In addition, participants *increase* their belief in the likelihood of the hypothesis, whilst they simultaneously *decrease* their belief in the reliability of the reporting source (P(H) to $P(H|Rep_1)$; *N* = 100), $BF_{10} = 5.958 * 10^7$. That is, the introduction of a dissenting minority source on belief in the likelihood of the hypothesis ($P(H|Rep_2) – P(H|Rep_3)$) leads to a significant decrease (*N* = 100), BF_{10} = 281255.7.

We next turn the high likelihood scenario (biologists predicting salmon growth). To test whether participants neither increase or decrease the reliability of sources that provide positive statements for highly likely hypotheses (hypothesis 2), we conducted a repeated measures ANOVA (P(Rel) to $P(Rel_1|Rep_1)$), finding no significant change (*N* = 100), BF_{10} = 0.401. We do however note the introduction of a contradicting source ($P(Rel_1|Rep_2) – P(Rel_1|Rep_3)$) leads to a significant decrease in reliability of the first source (*N* = 100), BF_{10} = 12458.36. Critically, given this introduction, and separating these results from those of the unlikely scenario, there was a substantial difference in the estimated reliability of the dissenter (μ = .489, σ = .207), and the first two (corroborating) reporters (Source 1: μ = .707, σ = .16; Source 2: μ = .717, σ = .156; *N* = 100), $BF_{10} > 1 * 10^{10}$. This suggests that - while introducing uncertainty to the reliability of the corroborating sources - providing dissenting reports about a hypothesis with a high prior belief and two corroborating reports can significantly damage perceived reliability. That is, if a minority dissents against prevailing wisdom *and* goes against other witnesses, she may suffer a loss of reliability.

We further note that as in the unlikely scenario, the introduction of a report from dissenting minority source on belief in the likelihood of the hypothesis ($P(H|Rep_2)$ – $P(H|Rep_3)$) leads to a significant decrease (*N* = 100), BF_{10} = $5.561 * 10^6$. The main results are shown in Fig. 3.

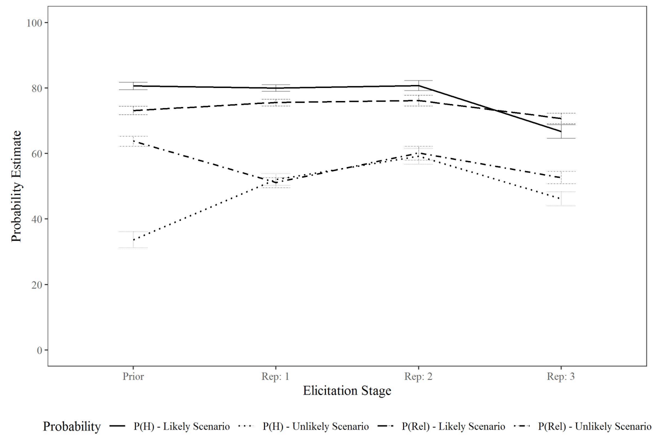

Fig. 3: P(Rel) and P(H) given $reports_{1-3}$

Results of shared reliability

To test whether the impact of introducing a shared reliability among sources (hypothesis 3), we compare posterior degrees of belief in the hypothesis and the reliability of the sources.

A repeated measures ANOVA was conducted on belief in the hypothesis (P(H)) for the introduction of the shared reliability information (i.e. $P(H|Rep_3)$ to P(H|SR)), with the inclusion of the shared reliability condition (high/low-quality) as a between-subjects condition.

For the unlikely scenario, belief in the hypothesis (economic crash), was affected by the introduction of a

[5] The high likelihood scenario was identical to the above, but considered predictions that the Norwegian salmon population would grow over the next 5-year period.

[6] All analyses assume an uninformed prior.

shared reliability (main effect of introduction), $BF_{Inclusion}$[7] = 110.1, and if shared reliability was high or low quality (low < high), $BF_{Inclusion}$ = 272.1, demonstrating a successful manipulation check. Importantly, the significant interaction of shared reliability condition, and its introduction, $BF_{Inclusion}$ = 550.8, revealed belief in the hypothesis decreased when the shared reliability was low-quality but increased when the shared reliability was of high quality. Consequently, the model with all the above terms included was the best fit, BF_M[8] = 550.81, and significant overall, $BF_{10} = 486.26$.

We observe the same effects for revision of reliability estimates. The main effect of an introduction of a shared reliability, $BF_{Inclusion} = 4.798 * 10^9$, and main effect of shared reliability condition (low-quality < high-quality), $BF_{Inclusion} = 1.583 * 10^8$, are best described by the significant interaction of the two, $BF_{Inclusion} = 4.421 * 10^8$, where high quality shared reliability leads to a minor increase in estimated reliability, whilst low quality shared reliability leads to a substantial decrease. Once again, the model with all the above terms included was the best fit, $BF_M = 4.421 * 10^8$, and significant overall, $BF_{10} = 1.552 * 10^{10}$.

The above analyses were then repeated for the likely scenario, where, against predictions, the belief in the hypothesis (salmon growth) was found to be unaffected by the introduction of a shared reliability, $BF_{Inclusion}$ = 1.227, or its quality, $BF_{Inclusion}$ = 0.194. However, the introduction of shared reliability was found to decrease estimations of source reliability, $BF_{Inclusion} = 1.116 * 10^{10}$, and whether a shared reliability was high or low-quality led to higher or lower reliability estimates (respectively), $BF_{Inclusion} = 5.082 * 10^7$, once more passing the manipulation check. Critically, reductions in reliability (given the introduction of a shared reliability among sources), is found to be localized to when the introduced shared reliability is of low-quality (right-hand facet, Fig. 4), $BF_{Inclusion} = 8.252 * 10^7$. Finally, the model with the above terms included was the best fit, $BF_M = 8.252 * 10^7$, and significant overall, $BF_{10} = 6.162 * 10^{10}$. The main results are shown in Fig. 4.

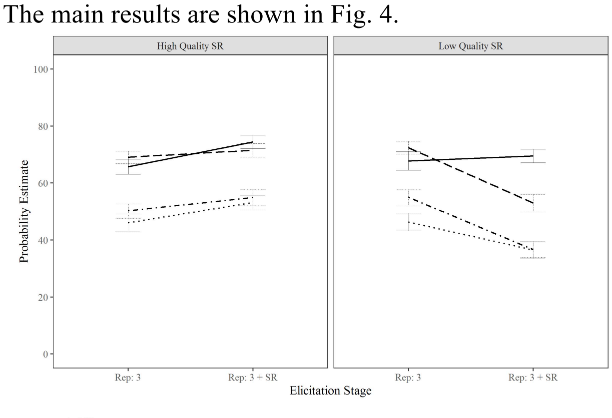

Fig. 4: P(Rel) and P(H) given shared reliability

[7] $BF_{Inclusion}$ shows the change in odds from the sum of the prior probabilities of models including the effect, to the sum of the posterior probabilities of models including the effect.

[8] BF_M shows the change from prior to posterior odds, given the model.

Discussion and concluding remarks

Despite the prevalence of dissent in public debates, the role of minority dissenters has not been adequately explored or modelled. This is a crucial function if to understand the functional impact of dissent in debates such as climate change or political predictions.

The paper tests how people revise beliefs in the reliability of sources and a hypothesis given sequential reports. The two initial reports support the hypothesis while the 3rd report rejects it. First, P(Rel) initially decreases when the source provides a positive report for an unlikely hypothesis, but rebounds when the 2nd source corroborates the initial report. Additionally, P(H) increases for the same reports while P(Rel) does not change for predicting the *likely* hypothesis while P(H) increases slightly. This replicates findings from Madsen et al (2018) and follows Bayesian predictions.

Second, the negation of the hypothesis yielded novel results. In both scenarios, negation decreased P(Rel) for all sources, presumably as it introduces noise and uncertainty. P(H) decreases when the 3rd source rejects the hypothesis for the unlikely scenario, but does not decrease for the likely scenario. This suggests that while participants revise their belief in an unlikely hypothesis (a market crash within six months), they decrease their belief in the hypothesis when dissent is voiced against this idea. Comparatively, P(H) does not decrease with dissent in the likely scenario. Rather, P(Rel) for $source_3$ decreases significantly given rejection of the likely hypothesis. P(Rel) also decreases for $sources_{1\text{-}2}$ given dissent in the likely case, but not to the same extent as is suffered by the dissenter.

Finally, shared reliability appears to work asymmetrically for consenters and dissenters. If the school enjoys a good reputation, perceived reliability increases for consenters, but less so for dissenters. If the shared reliability is perceived as high quality, people's degree of belief in the hypothesis additionally increases. However, for both perceived source reliability and the hypothesis, we see a decrease when the shared reliability is of poor quality.

In all, the study suggests people are sensitive to source reliability as well as the structural relationship between the sources. Belief revisions generally follow Hahn et al. (2009) such that positive reports from very credible sources lend credence to the hypothesis. Additionally, participants update perceived source reliability in accordance with predictions, as supporting unlikely hypotheses is initially detrimental, but sequentially rebounds given corroboration. Finally, perceived partial dependence is crucial, as shared reliability moderates perceptions of the hypothesis and the reliability of all sources involved. In all, the study provides additional support for a Bayesian approach to source reliability.

References

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

Chaiken, S. & Maheswaran, D. (1994) Heuristic Processing Can Bias Systematic Processing: Effects of Source Credibility, Argument Ambiguity, and Task Importance on Attitude Judgement, *Journal of Personality and Social Psychology* 66 (3), 460-473

Cialdini, R. B. (2007) *Influence: The Psychology of Persuasion*, Collins Business

Cuddy, A. J. C., Glick, P. & Beninger, A. (2011) The dynamics of warmth and competence judgments, and their outcomes in organizations, *Research in Organizational Behavior* 31, 73-98

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

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

Hahn, U., & Oaksford, M. (2006) A normative theory of argument strength, *Informal Logic* 26*,* 1-24

Hahn, U., & Oaksford, M. (2007) The rationality of informal argumentation: A Bayesian approach to reasoning fallacies, *Psychological Review 114*, 704-732

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*,* 1496-1533

Harris, P. L., & Corriveau, K. H. (2011). Young children's selective trust in informants, *Philosophical Transactions of the Royal Society B, 366,* 1179-1187

Howson, C., & Urbach, P. (1996). *Scientific Reasoning: The Bayesian Approach (2nd Edition).* Chicago, IL: Open Court

JASP Team (2018). JASP (Version 0.9)[Computer software].

Lagnado, D. A., Fenton, N., & Neil, M. (2013). Legal idioms: a framework for evidential reasoning, *Argument & Computation* 4 (1), 46-63.

Madsen, J. K. (2016) Trump supported it?! A Bayesian source credibility model applied to appeals to specific American presidential candidates' opinions, Papafragou, A., Grodner, D., Mirman, D., & Trueswell, J.C. (Eds.) *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, Austin, TX: Cognitive Science Society, 165-170

Madsen, J. K., Hahn, U. & Pilditch, T. (2018) Partial source dependence and reliability revision: the impact of shared backgrounds*, Proceedings of the 40th Annual Conference of the Cognitive Science Society*

Oaksford, M. & Chater, N. (1991) Against logicist cognitive science, *Mind and Language* 6, pp. 1-38

Oaksford, M. & Chater, N. (2007) *Bayesian Rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press

Perfors, A., Navarro, D. J. & Shafto, P. (2018) Stronger evidence isn't always better:A role for social inference in evidence selection and interpretation, *Proceedings of the 40th Annual Conference of the Cognitive Science Society*

Petty, R. E. & Cacioppo, J. T. (1984) Source Factors and the Elaboration Likelihood Model of Persuasion*, Advances in Consumer Research* 11, 668-672

Tormala, Z. L. & Jackson, J. J. (2007) Assimilation and Contrast in Persuasion: The Effects of Source Credibility in Multiple Message Situations, *Personality and Social Psychology Bulletin* 33 (4), 559-571

Whalen, A., Griffiths, T. L. & Buchsbaum, D. (2013) Sensitivity to shared information in social learning*, Cognitive Science* 42, 168-187