

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Partial source dependence and reliability revision: the impact of shared backgrounds

Permalink

<https://escholarship.org/uc/item/8t96w1mk>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Madsen, Jens Koed

Hahn, Ulrike

Pilditch, Toby D

Publication Date

2018

Partial source dependence and reliability revision: the impact of shared backgrounds

Jens Koed Madsen (jens.madsen@ouce.ox.ac.uk)

University of Oxford, School of Geography and the Environment
South Parks Road, OX1 3QY, Oxford, United Kingdom

Ulrike Hahn (u.hahn@bbk.ac.uk)

Department of Psychological Sciences, Birkbeck, University of London,
Malet Street, London, WC1E 7HX, U.K.

Toby D. Pilditch (t.pilditch@ucl.ac.uk)

Department of Experimental Psychology, University College London,
26 Bedford Way, London, WC1H 0AP, U.K.

Abstract

The paper explores how people revise their belief in a hypothesis and the reliability of sources given independence of sources or partial dependence (e.g. the sources share a background). Specifically, we test a formal model of reliability revision.

The study provides support for Bovens and Hartmann's model of reliability revision. If a source provides a positive report for an unlikely hypothesis, participants initially revise the reliability of the source negatively. However, as additional positive reports emerge, they increase their estimate of the reliability of the source. Further, if it becomes known the sources are partially dependent (here, taught the same school of thought), the reliability of the source decreases again. Both of these findings are in line with the Bayesian approach to reliability revision.

Keywords: Bayesian reasoning; Reliability revision; Sequential testimonies

Introduction

In everyday life, we constantly get information from sources such as meteorologists, friends, and co-workers. Further, multiple sources often provide reports about the same topic – for example, we may read different columns that forecast the economic impact of Brexit. Sources may be independent of one another (producing data and conclusions in isolation from other sources) or dependent (e.g. sources discuss prior to their individual reports or share a common background). The structure of dependence is crucial.

A failure to appreciate the dependence of information can lead to potentially disastrous conclusions. If an intelligence agency gets multiple reports concerning weapons of mass destruction in a foreign country, they may increase their belief in the veracity of this proposition (unlikely as it may be prior to the reports). Multiple cooperating reports may sway the agency to believe an improbable hypothesis. If it subsequently becomes known that all reports came from sources with a common, flawed approach, the cooperation of the reports is compromised. That is, a failure to appreciate the dependency or independence of sources is critical to reasoning and decision-making.

The paper tests belief revision processes concerning the hypothesis and reliability given source independence or partial dependence. More specifically, we take point of departure in a formal account of dependence and reliability (Bovens & Hartmann, 2003, chapter 3) and test their intuitions and predictions empirically.

The impact of source reliability on belief revision¹

The reliability of the source is crucial for reasoning and decision-making and requires formal modeling. It influences a range of human cognitive phenomena such as the reception of persuasive messages (Petty & Cacioppo, 1984; Tormala & Clarkson, 2007), the development of children's perception of the world (Harris & Corriveau, 2011), legal reasoning (Lagnado et al., 2013), adherence with persuasion strategies (Cialdini, 2007), and how people are seen in social situations (Fiske et al., 2007; Cuddy et al., 2011).

Cognitive and social psychology has approached this source reliability question in a number of different ways. Reliance on the reliability of others has been described as a shallow persuasive cue (Petty & Cacioppo, 1984), and, alternatively, as rationally justified (Hahn et al., 2009).

Harris et al (2015) tested a model that amalgamates two components of reliability: perceived trustworthiness and perceived expertise. Expertise refers to the *capacity* of providing accurate information about the topic. This is domain-dependent. For example, a carpenter may provide relevant and accurate information about types of wood and how they should be handled, but may not be able to provide guidance in neurosurgery. Trust, on the other hand, refers to the *intention* of providing true and accurate information to the best of their ability. For example, if the carpenter has a motive to sell surplus wood, she may falsely claim a particular type of wood is useful even in situations where it is not. These components are orthogonal such that a person can be highly expert, but very untrustworthy and vice versa.

¹ Some papers describe source credibility and others reliability. To ease the reading experience, we only use 'reliability'.

Bayesian approaches to reasoning take point of departure in subjective, probabilistic degrees of beliefs in propositions where the posterior degree of belief in a proposition is captured by Bayes' theorem (Oaksford & Chater, 2007; Howson & Urbach, 1996). The approach is suggested as an alternative to logicist approaches to reasoning (Oaksford & Chater, 1991) and has been applied to argumentation theory (Hahn & Oaksford, 2006; 2007). The findings suggest Bayesian reasoning can account human information integration in practical reasoning.

In Harris et al (2015), Bayes' theorem integrates elements of reliability such that the posterior degree of belief in the hypothesis (H) given the representation (Rep) yields:

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

This allows agents to use their perception of the reliability (trustworthiness and expertise) of the source to update their belief given a positive or negative report from that source.

This framework has been tested empirically and has been shown to capture how people update their beliefs given information from more or less reliable independent sources (Harris et al., 2015; Madsen, 2016).

Super reliability and reliability revision

Studies of belief revision and propagation tend to focus on subjects' belief in the hypothesis, either as a result of new evidence (e.g. Bayesian belief revision models of reasoning and argumentation) or as a result of reports from other sources (e.g. Bayesian source reliability models). This focus makes sense, as the belief in the hypothesis motivates or influences decisions such as voting, economic behavior, etc. while the reliability of the source can be seen as an auxiliary factor in updating the belief in the hypothesis.

These studies indicate that perceived reliability influences belief revision concerning the hypothesis. This means the perception of source reliability *itself* is important in the belief revision process. Therefore, if the perceived reliability of the source changes, Bayesian normative models suggest the subsequent impact of that source should change also. For example, the Boy Who Cried Wolf made continuous bad forecasts (willingly), which caused the villages to decrease their reliability estimate of the boy (with disastrous consequences).

Bovens and Hartmann (2003) provide a foundational perspective on modeling source reliability from a Bayesian perspective. In their book, they show the structure and relationship of sources influence the degree to which their reports should normatively shape the recipient's belief in the hypothesis *and* the posterior degree of belief in the reliability of each source given multiple testimonies. Sources may be independent of one another (fig. 1) or be linked in a partially dependent manner (fig. 2).




Fig. 1: Independent condition. Each source has an independent reliability (Rel) and provides a report (Rep) conditionally independent of other sources.

This condition refers to instances where the sources are considered entirely independent of one another. For example, climate scientists may conduct independent studies of the same phenomenon and produce reports of their findings without any knowledge of the conclusions of the other teams. This would constitute fully independent sources, as they do not rely on the same apparatus, do not share results before making their reports known, and do not communicate between teams. However, sources may share a common background (e.g. attending the same school). If so, the sources become partially dependent (see fig. 2 for an example of dependence).




Fig. 2: Super Reliability condition. All source reliabilities now share a common ancestor "super reliability" (e.g. shared background).

In fig. 2, the sources share a background. This provides a constraint on the informativeness of each source, as their reliabilities are influenced by a common-cause (e.g. economists coming from the same good or bad school). Assuming interpretation of the source does not allude to a myriad of conclusions, if the reliability of sources are now conditional on a common background, that background can weaken the impact of the reports provided by those sources. For example, if several doctors provide diagnoses for a

² $P(Rep|H) = P(Rep|H, exp, T) * P(Exp) * P(T) + P(Rep|H, \neg exp, T) * P(\neg Exp) * P(T) + P(Rep|H, \neg exp, \neg T) * P(\neg Exp) * P(\neg T) + P(Rep|H, exp, \neg T) * P(Exp) * P(\neg T)$; mutatis mutandis for $P(Rep|\neg H)$

patient, it makes an operational difference to the impact of their reports if they were found to all have attended the same low standard medical course. In comparison to a fully independent case, recipients should treat reports from these doctors as partially compromised. This shared background not only influences the reliability of each source, but in turn influences the degree to which reports from those sources impact the hypothesis.

Bovens and Hartmann (2003) provide a formal way to calculate "...how the posterior probability of the reliability of the n^{th} witness increases as more and more witness reports from in" (p. 79):

$$P^{*(n)}(\text{REL}_n) = P(\text{REL}_n | \text{REL}_1, \dots, \text{REL}_n) \\ = \frac{h[us(s+a\bar{s})^{n-1} + \bar{u}t(t+a\bar{t})^{n-1}]}{h[u(s+a\bar{s})^n + \bar{u}(t+a\bar{t})^n] + \bar{h}[u(s+a\bar{s})^n + \bar{u}(t+a\bar{t})^n]}$$

where u refers to the probability of reliability of the SR, $P(\text{SR})$, s refers to the conditional probability: $1 > p(\text{REL}_i | \text{SR})$, and t refers to the conditional probability $> p(\text{REL}_i | \text{SR}) > 0$, a refers to a randomization parameter, and h refers to the prior probability of the hypothesis³.

They show 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 to .25 (see p. 80), but increases as additional positive reports confirm the initial report. Importantly, the current study does not elicit a randomization parameter specifically. Consequently, we test if the overall intuitions are in line with observed data. We describe the hypotheses in more detail in the following.

Present Research

In the paper, we explore four hypotheses. First, we explore whether participants revise their belief in the reliability of the source in line with Bovens and Hartmann (2003). For improbable hypotheses (in chapter 3, Bovens and Hartmann cite $p(h) = 0.3$), a single positive report should decrease the reliability of the source. However, as independent sources continue to cooperate and supply positive reports, the reliability of the sources should increase. That is, $p(\text{rel})$ should initially decrease given unlikely statements, but then increase with additional reports.

Second, in line with Bovens and Hartmann, sources that provide positive statements for highly likely hypotheses (e.g. $p(h) = .9$) should neither increase nor decrease, as they are merely confirming what is a priori very likely.

Third, we explore whether source independence adjusts reliability estimations. If sources are wholly independent, they should, normatively, update in line with the predictions tested in the first hypothesis. If, however, sources are dependent, this pattern should change. Dependency may result from direct communication between the sources (e.g. members of a jury will eventually provide identical verdicts

of guilt, as the function of the jury is to reach a consensus beyond reasonable doubt), from a shared source (sources may provide identical information if they have all consulted the same data set or book), or via sharing a common reliability ancestor (such as working in the same organization or having attended the same schooling). The latter is a case of Super Reliability (SR).

Direct communication and source consensus implies strong dependence, as the sources are required to share their beliefs (e.g. a jury). Comparatively, a shared background implies weaker dependence, as individual source reliabilities mediate the impact of the overarching super reliability (e.g. having attended the same formal training still allows for differences in how well individuals use it). While both types of dependency are certainly of interest, the current paper explores dependency through a shared background. This is in line with forecasting literature (see e.g. Hogarth, 1989; Soll, 1999)

Finally, we explore whether participants adjust reliability estimates of sources retrospectively, or if additional reports only reflect on the most recent sources. That is, after the first report, participants provide their reliability estimate for the first source (the one that provides the positive report). We explore whether seeing subsequent positive reports for the same hypothesis leads to a revision of the reliability of the original source despite the fact that this source does not contribute with additional reports.

If participants revise their beliefs about the reliability of the source retroactively, we should see no differences between estimates of source reliability given new reports, as previous sources are also revised in light of new reports. If, however, participants do not revise beliefs retrospectively, the reliability of sources should differ, as participants learn additional information.

Method

To test the above hypotheses, we employ the following methodological manipulations: To test H1 & H2, the prior probability of the hypothesis is manipulated as high/low. This allows exploration of whether reliability initially decreases and then subsequently increases given additional positive reports for highly unlikely statements (H1) and if providing positive reports for highly probably does not exhibit this effect (H2). To test H3, additional corroborating sources are incrementally introduced, followed by an SR manipulation. SR was presented as either high reliability (a school with an excellent reputation) or low (a school with a poor reputation). This explores the sensitivity to the goodness of the SR. To test H4, $p(h)$ as well as $P(\text{rel})$ estimates are elicited after each report for the hypothesis as well as for every source (meaning $p(\text{rel})$ of source₁ is elicited thrice, once after each positive report).

Materials and procedure

Materials: In order to test the above hypothesis, two scenarios were used. In scenario 1 (low probability condition), participants were asked to consider the

³ See appendix C.3 and C.4 in Bovens and Hartmann for the full derivation.

likelihood of a market crash within a 6-month period. Specifically, they see the following:

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?

Scenario 2 (high probability condition) considers the likelihood that the salmon population will grow within a 5-year period. Specifically, they see the following:

Imagine you are watching a nature programme about fish. Specifically, the programme considers whether or not the salmon population of Norway will grow (i.e. increase by more than 10%) over the next 5 years. Historically, the likelihood of an increase in the salmon population in Norway within a 5-year window is 85%.

In your opinion, how likely is the salmon population of Norway will grow over the next 5 years?

In addition to the scenarios, participants were presented with statements from experts in the field (economist and biologists). This allowed for reliability measures of the sources. For the biological scenario, they saw the following:

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 biologists in predicting the growth of species?

To generate reports about the hypothesis, participants were told experts had been interviewed. Specifically, they saw the following:

Now, imagine that a biologist, Linda, is being interviewed about the salmon. Linda states the following: "I am completely certain the salmon population of Norway will grow over the next 5 years."

Given Linda's report, how likely is the salmon population of Norway will grow over the next 5 years?

Finally, to generate SR conditions (high and low), the participants were told the experts had attended the same school. Specifically, they saw the following:

It turns out, all the interviewed biologists studied at the same school and subscribe to the same biological models. Their school has a very good reputation for excellent teaching and accurate approaches to biology [HIGH SR condition]// Their school has a very bad reputation for sloppy teaching and out-dated approaches to biology [LOW SR condition].

Given the fact that they all studied at the same school and follow the same biological models, how likely is the salmon population of Norway will grow over the next 5 years?

In all, the material includes prior beliefs, $p(h)$ and $p(\text{rel})$, reports, posterior beliefs, $p(h|\text{rep})$ and $p(\text{rel}|\text{rep})$ as well as SR manipulation, high and low.

Procedure: Participants first provide 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) and their belief in the reliability of the type of source (economist or biologist) from 0-1 (0: biologists are completely unreliable; 1: biologists are completely reliable).

Having provided their priors, participants saw sequential reports from experts (in total, participants read 3 reports, all of whom were positive). After each report, participants provided their degree of belief in the hypothesis as well as their degree of belief in the reliability of each source (the reliability source₁ was elicited thrice, but the reliability of source₃ was only elicited once after the third report).

Finally, participants were asked to "...consider two possible continuations to the scenario, providing your assessments for each". They then read both SR conditions and were asked to provide their degree of belief in the hypothesis and in the reliability of each expert given the dependency between the experts.

The study was a within-subjects design to decrease noise from interpretation. They responded to all questions for each scenario before seeing the 2nd scenario. All conditions were counterbalanced (50% read the market crash first; 50% saw high SR first). Thus, participants read two scenarios (low and high probability), read 3 reports for each scenario and read two hypothetical continuations of the scenario (high and low SR conditions).

Participants

100 participants (68 female, $\mu_{\text{age}} = 31.88$, $\sigma = 10.34$) 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%.

Average completion time was 5.63 min ($\sigma = 2.95$) and participants were paid £1.10 (resulting in an effective fair hourly wage of roughly £11.7/hour for participation). No participants were excluded from the analysis.

Results

The probability manipulations were successful in generating high and low estimates for the two scenarios: The market crash scenario was rated as unlikely ($\mu = .29$, $\sigma = .23$) and the salmon growth scenario was rated as likely ($\mu = .77$, $\sigma = .15$). The unlikely scenario was particularly fortunate, as the example from Bovens and Hartmann uses $h = .3$, making the current scenario comparable to their example. In Bovens and Hartmann, $p(\text{rel}) = .5$. Both sources in our scenarios were rated higher ($p(\text{rel_economist}): \mu = .64$, $\sigma = .18$; $p(\text{rel_biologist}): \mu = .73$, $\sigma = .15$). Importantly, though, both sources were rated positively, which allows to test whether positive reports of unlikely hypotheses influence reliability estimates negatively.

To test whether participants revise their belief in the reliability of the source in line with Bovens and Hartmann (hypothesis 1), we explore if participants adjusted reliability estimates in the initial source given sequential testimonies.

In chapter 3, Bovens and Hartmann provide a model that shows positive reports of an unlikely hypothesis should initially decrease estimates of reliability.




Fig. 3: $p(h)$, $p(h|rep_{1-3})$, $p(rel_source_1)$, $p(rel_source_1|rep_{1-3})$ for both conditions ($p(h) = .29$ and $p(h) = .77$)

In line with the predictions from Bovens and Hartmann, we observe a negative revision of reliability of source 1 given a positive report of an unlikely hypothesis (in the current design, the source predicts the stock market will crash within a 6-month period). However, as other participants learn other experts provide identical reports, they revise their belief in the initial source and revise reliability in a positive direction. This specifically tested a scenario with a low prior probability (here, $p(h) = .29$).

In addition, participants *increase* belief in the likelihood of the hypothesis whilst they simultaneously *decrease* belief in the source reliability. While the paper tests the former (change in reliability), it is worth noting the participants follow previous findings (e.g. Harris et al., 2015) that show $p(h|rep)$ increases if the source is viewed as reliable.

To test whether sources that provide positive statements for highly likely hypotheses neither increase nor decrease their reliability (hypothesis 2), we conduct a one-way ANOVA to test if the reports change the estimation of the reliability of the biologist. While there is a slight increase in reliability across conditions ($\mu_{prior} = .73$, $\sigma = .15$; $\mu_{first} = .75$, $\sigma = .16$; $\mu_{second} = .76$, $\sigma = .15$; $\mu_{third} = .78$, $\sigma = .15$), we see no significant difference between conditions ($df = 3$, $F = 2.061$, $p = .105$). This suggests the reliability estimates of reliable sources providing positive reports for likely statements neither increase nor decrease. This can be seen visually on fig. 2 where $p(rel)$ remains fairly flat.

To test whether source independence adjusts reliability estimations (hypothesis 3), we compare posterior degrees of belief in the hypothesis and the reliability of the sources. As described in the above, we manipulate the reliability of the SR and compare $p(h|rep_3)$ with $p(h|SR)$ and $p(rel|rep_3)$ with $p(rel|SR)$ for high and low SR conditions.

In line with predictions, a one-way ANOVA shows significant decreases the degree of belief in the hypothesis when sources are dependent for the unlikely scenario (crash:

$\mu_{prior} = .65$, $\sigma = .24$, $\mu_{bad_SR} = .32$, $\sigma = .23$, $\mu_{good_SR} = .55$, $\sigma = .22$, $df = 1$, $F = 32.91$, $p < 0.001$). For the likely scenario, the effect is less strong (salmon growth: $\mu_{prior} = .81$, $\sigma = .14$, $\mu_{bad_SR} = .64$, $\sigma = .20$, $\mu_{good_SR} = .78$, $\sigma = .15$, $df = 1$, $F = 19.66$, $p < 0.001$). We observe the same tendency for degree of belief in the reliability of the sources (salmon growth: $\mu_{prior} = .78$, $\sigma = .15$, $\mu_{bad_SR} = .50$, $\sigma = .20$, $\mu_{good_SR} = .77$, $\sigma = .16$, $df = 1$, $F = 33.95$, $p < 0.001$; crash: $\mu_{prior} = .62$, $\sigma = .22$, $\mu_{bad_SR} = .34$, $\sigma = .23$, $\mu_{good_SR} = .56$, $\sigma = .22$, $df = 1$, $F = 27.73$, $p < 0.001$). In addition, the SR manipulation works, as paired-sample t-tests show low-quality dependency is significantly lower than high-quality dependency. This holds for belief in the hypothesis (crash: $t = 5.704$, $p < 0.001$; salmon: $t = 6.735$, $p < 0.001$) and reliability (economist: $t = 5.782$, $p < 0.001$; biologist: $t = 9.763$, $p < 0.001$). For a visual overview, see fig. 4.




Fig. 4: $p(h)$, $p(h|rep_{1-3})$, $p(rel_source_1)$, $p(rel_source_1|rep_{1-3})$, $p(h|SR_bad)$, $p(h|SR_good)$, $p(rel|SR_bad)$, and $p(rel|SR_good)$ for both conditions

To test whether participants adjust reliability estimates of sources retrospectively, or if additional reports only reflect on the most recent sources (hypothesis 4), we conduct a one-way ANOVA between the source reliability estimates of all sources. Specifically, we compare participants' estimates of the reliability of source₁ after all reports have been provided. The ANOVA shows no difference between conditions ($df = 2$, $F = .060$, $p = .942$). This suggests people are update their belief in source₁ retroactively given new reports from additional sources even if source₁ does not contribute with additional reports.

Discussion and concluding remarks

The paper explores how sequential testimonies and partial dependence modulates reliability estimates of sources. We explored four hypotheses:

First, we tested whether participants revised their posterior degree of belief in the reliability of the source in line with Bovens and Hartmann (2003). The data supports this prediction, as $p(rel)$ initially decreased given a positive

report of an unlikely hypothesis ($h = .29$), but subsequently increased as more positive reports came through.

Second, we tested whether sources that provide positive statements for highly likely hypotheses neither increase nor decrease their reliability. The data provides indicative support for this, as reliability of sources remained constant in the scenario with a likely prior probability ($h = .77$).

Third, we explored whether source independence adjusted reliability estimations. The data provides support for this hypothesis. When participants learned the experts attended the same school, they adjusted their posterior degree of belief negatively, both for the hypothesis and the reliability of each source. In line with expectations, the effect was stronger if the experts' school was bad compared with the scenario where the school was described as excellent.

Finally, we test if participants revised posterior degree of belief in the reliability of the source retrospectively. The data supports this hypothesis, as source₁ initially decreased when reporting an unlikely hypothesis. Yet, as sources_{2,3} provided similar reports, reliability of source₁ was adjusted in line with the n^{th} source (enjoying the same reliability as source₂ after 2 reports, *mutatis mutandis* with source_{2,3} after 3 reports). Overall, the data provides preliminary support for the model provided by Bovens and Hartmann.

Future work

We stress the exploratory nature of the study, as we did not elicit a specific randomization parameter and only tested two scenarios. One of these were specifically designed with a high prior probability in the hypothesis, as this should not reduce or increase the reliability of the source. As such, this case functions as a control scenario. Additionally, we did not explore a range of different dependency structures such as direct communication and consensus (the jury case), one-way dependency (source n can see the reports of source $n-1$ before making her statement whilst source $n-1$ cannot see the reports of subsequent sources), or other types of information structures between sources. Future work should test reliability updating given a much wider range of social and information structures, a wider range of hypotheses, different signal strength, and differences in SR.

The study offers confirmatory evidence for the predictions made by Bovens and Hartmann. Formalising the dynamics of reliability and hypothesis revision is of considerable interest for dynamic systems exploring belief diffusion or propagation in social networks (e.g. Vallinder & Olsson, 2014; Duggins, 2016).

References

- Bovens, L., & Hartmann, S. (2003). *Bayesian epistemology*. Oxford: Oxford University Press.
- 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
- Duggins, P. (2016) A Psychologically-Motivated Model of Opinion Chance with Applications to American Politics, *arXiv:1406.7770*
- 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
- Hogarth, R. M. (1989) On combining diagnostic "forecasts": Thoughts and some evidence. *International Journal of Forecasting* 5, 593-597.
- Howson, C., & Urbach, P. (1996). *Scientific Reasoning: The Bayesian Approach (2nd Edition)*. Chicago, IL: Open Court
- 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
- 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
- Petty, R. E. & Cacioppo, J. T. (1984) Source Factors and the Elaboration Likelihood Model of Persuasion, *Advances in Consumer Research* 11, 668-672
- Soll, J. B. (1999) Intuitive theories of information: Beliefs about the value of redundancy. *Cognitive Psychology* 38 (2), 317-346
- 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
- Vallinder, A. & Olsson, E. J. (2014) Trust and the value of overconfidence: A Bayesian perspective on social network communication, *Synthese* 191, 1991-2007