Synthese (2025) 205:143
https://doi.org/10.1007/s11229-025-04969-w

ORIGINAL RESEARCH

The problem of dependency

Toby D. Pilditch[1,2] · Ulrike Hahn[3] · David Lagnado[1]

Received: 25 May 2024 / Accepted: 19 February 2025

Abstract

It is a wide-spread assumption that multiple pieces of evidence, whether they involve measurements or testimony, provide stronger evidential support when they are independent than when they are not. The standard view is that non-independence creates redundancy and leads to over-confidence (see e.g., Soll in Cogn Psychol 38(2):317–346, 1999; Nisbett and Ross in Human inference: strategies and shortcomings of social judgment. Prentice Hall, New Jersey, 1980). In keeping with this, fields as diverse as law, statistics, and philosophy have sought to understand the implications of non-independence and set out rules for dealing with it. In this paper, we set out how the challenges posed by non-independence are both far more wide-spread and more challenging than typically assumed. Specifically, we review how Bayesian probabilistic analysis reveals "simple" inference cases that conflict with wide-spread assumptions about the value of independence. These suggest that mere intuition is not a reliable guide in this arena. However, we then show that the same framework cannot (in its current form) be scaled to common real-world situations even in principle. We conclude with a discussion of the more modest, partial, solutions that are currently available and outline possible future directions.

Keywords Dependencies · Epistemics · Reasoning under uncertainty · Testimony

✉ Toby D. Pilditch
t.pilditch@ucl.ac.uk

1 Department of Experimental Psychology, University College London, 26 Bedford Way, London WC1H 0AP, UK

2 School of Geography and the Environment, University of Oxford, South Parks Road, Oxford OX1 3QY, UK

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

Published online: 28 March 2025

1 Dependence and independence

Imagine trying to judge whether there will be a thunderstorm tomorrow. You already have an initial sense of how likely this is, but seek out further information. Your friend Alex says that he thinks there will be a thunderstorm. You believe Alex to be trustworthy (he's your friend!), but he's not an expert meteorologist. You then bump into a second friend, Jessie, who also thinks there will be one. This corroboration, intuitively, further bolsters your belief in a thunderstorm tomorrow.

Already, there are a number of relevant factors to consider in your evaluation. Are thunderstorms common in your area? How difficult is it for lay people to accurately predict thunderstorms a day in advance, so what weight should you give to their judgments? Such questions must be addressed (implicitly or explicitly) whenever we adjust our views in response to those of others. However, there is an important further issue here: how are Alex' and Jessie's judgments related to each other? Does it matter whether they each came to their own separate conclusion or just both heard the same radio forecast? If they both had their own reasons *and* heard the same forecast, is there some way to avoid 'double counting' the common forecast?

This is the central problem posed by the non-independence of evidence sources. It applies not just to testimony (Alex and Jessie) but also to measurement devices. For example, should we prefer a second measurement with the same device or should we prefer to receive a measurement from a second, different, device? The question of whether to prefer repeated measurements or a second device has seen decades of discussion in the philosophy of science and epistemology under the header of the 'variety of evidence thesis' (e.g., Fitelson, 1996) and impacts everything from choice of physical devices to discussions about the value of replication within psychology (Oberauer & Lewandowsky, 2019). It is, at heart, a question about the value of independence. At the same time, there are many practical contexts where some degree of non-independence is simply a fact one can do little about: current climate models, for example, share components raising questions about how to account for this fact in evaluating ensemble predictions; scientists share education, training and methods; court cases are stuck with the actual witnesses that could be found.

Working out how we should respond to non-independence is a normative question, that is, a question about how we *ought* to behave. In response, there have been many proposals of both informal and formal rules—in legal systems, science, and everyday life. Ideally, prescriptions of how we *should* form beliefs should rest on more than mere intuition or say so, but rather have some independent justification. In this paper, we detail both what marks out the Bayesian framework as a suitable tool for addressing questions of dependence and some of the insights its applications have afforded. Those applications, in effect, treat people and testimony just like any other evidence source such as physical measurement devices. The main contribution of this paper, however, is to make clear the limits of such an approach. Specifically, we show how and why the current Bayesian tool kit, in particular Bayesian Belief Networks (BNs), fails to deal with the fundamental fact that much human exchange involves communication across social networks. In other words, even where we ostensibly exchange information in pairwise exchanges (such as with Alex and Jessie), we are part of a wider network of social exchange because (some of) those we speak to also speak to each other. We are, in effect, individuals in collectives (see Hahn, 2023).

Given that violations of independence may lead to over-weighting of evidence, and that there is considerable concern about the way online social networks, in particular, may fuel the development of extreme opinions, it is revealing that the same Bayesian tools that work well with a small number of sources, do not work in this more general case. Furthermore, the reasons for those are not just practical or computational considerations, but rather fundamental representational limitations. This, in turn, shows the challenges posed by non-independence to be deeper than previously assumed.

Specifically, the paper proceeds in two main parts. We first briefly detail ways in which the Bayesian framework is seen as particularly suited for answering the normative questions posed by non-independence and review some of its success here. We then outline how and why the same tools fail once we extend sources to even a small social network. We conclude with thoughts on future directions.

2 Bayesian tools for the normative question

In this section, we introduce the relevant aspects of the Bayesian framework, illustrate why its normative guidance for capturing dependence when combining multiple testimonies (Sect. 2.1), and describe how the framework has contributed to elucidating the variety of evidence thesis (Sect. 2.2). We then distinguish different types of dependence (Sect. 2.3). Finally, in Sect. 2.4, we draw out those features of the framework that enable these successes in preparation to showing, in the remainder of the paper, how those features become crucial stumbling blocks for other cases of dependence.

Arguably, the Bayesian framework is sufficiently established that no further words on its suitability for conceptualising normative questions are required. It not only has a comparatively well-developed normative foundation (Hayek, 2008; Pettigrew, 2016), Bayesian models of cognition exist across many domains including learning (e.g., Bramley et al., 2017; Kruschke, 2006), reasoning (e.g., Oaksford & Chater, 1994; Sloman & Lagnado, 2015), argumentation (e.g., Hahn & Oaksford, 2007), and prediction (e.g., Griffiths & Tenenbaum, 2006). Bayesian approaches have also seen a resurgence in statistics (see e.g., Wagenmakers, 2007), gaining territory on more traditional, frequentist statistical techniques. Most importantly, the Bayesian framework has been widely used in both the philosophy of science and in epistemology in order to elucidate questions about how humans can arrive at secure knowledge (e.g., Bovens & Hartmann, 2003).

We think it is nevertheless useful to be more explicit about the reasons for focusing on the Bayesian framework in the present context. Our paper seeks to identify both circumstances where a Bayesian framework clarifies the problem of non-independence, and, more importantly, cases where it cannot. That even an ideal agent model cannot (currently) adequately capture these circumstances is, we think, informative with respect to understanding the full scale of the challenge human cognitive agents are trying to solve and it helps see clearly the difficulties for finding practical solutions. In short, for Bayesian sceptics, the framework may very much be treated in this paper as a crutch that is ultimately thrown away.

There have been many informal, qualitative rules proposed for dealing with witness testimony, such as legal rules on hearsay or rules proposed in the argumentation literature (see Walton, 2007). However, unless one wants to simply discard any type of evidence with dependence (a step which is arguably neither desirable nor possible in many circumstances as will become clear), the normative question itself calls for a *quantitative answer*. It may be intuitive that dependence is a kind of redundance as widely assumed (e.g., Nisbett & Ross, 1980), but that alone doesn't answer by *how much* we should adjust our conclusions in response. The Bayesian framework promises such quantitative answers.

The final, and possibly most important, reason for adopting the Bayesian framework for the current investigation, however, is that the problem of non-independence has long been one of its central concerns. Much of the scepticism toward probability theory within computer science and artificial intelligence in the 1970's and 80's arose specifically from the fact the problems non-independence poses for the application of Bayes' rule in practice. The combinatorial explosion created by multiple non-independent variables, however made many Bayesian calculations unworkable. Historically, the key insight was realising that conditional independence assumptions could be exploited to reduce that complexity (Dawid, 1979, 1980; Spohn, 1980). Hence, the major turning point in the theoretical and practical application of Bayesian statistics was the development of so-called Bayesian Belief Networks (BNs) (see e.g., Pearl, 1988). BNs provide a graphical representation of probabilistic independence relationships between variables which allows one to identify large parts of the joint probability distribution as irrelevant, meaning those probabilities no longer need to be specified. Figure 1 contains a simple three-variable example.

Sprinkler	
On	.4
Off	.6

Rain	
True	.1
False	.9

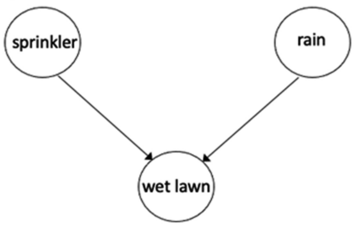

Wet Lawn	Sprinkler = On Rain = True	Sprinkler = On Rain = False	Sprinkler = Off Rain = True	Sprinkler = Off Rain = False
True	.99	.8	.6	.01
False	.01	.2	.4	.99

Fig. 1 Simple three variable Bayesian Belief Network, with underlying probability tables. *Note*: Nodes represent variables, and the directed arrows represent dependence relationships. The three tables are the conditional probability tables for the variables 'sprinkler', 'rain', and 'lawn' respectively

Though this will be familiar territory to most readers, it will help with subsequent distinctions to nevertheless go over this trivial example. The simple model of Fig. 1 represents the fact that whether or not the lawn is wet depends both on the state of the sprinkler (on/off) and whether or not it is raining. As a consequence, calculating the probability of the lawn being wet (or, for example, inferring the probability that it is raining, given that we have observed a wet lawn) requires one to know the probabilities for four distinct cases: sprinkler on & rain; sprinkler on & no rain; sprinkler of & rain; sprinkler off & no rain. These probabilities must be specified in a so-called conditional probability table (CPT) in order to perform the desired computations. At the same time, however, the model of Fig. 1 tells us that there are a number of joint probabilities we do *not* need to know. Specifically, there is no (direct) link between 'sprinkler' and 'rain' which indicates that (in this particular back yard, at least) whether or not the sprinkler is on or off is independent of whether it is raining (e.g., because it is simply on a timer). As a result, we do not need to specify (or even know) how likely it is that the sprinkler will be on at the same time that it is raining. We only need to know how likely it is that the sprinkler is on or off, and how likely it is that it is or is not raining. As a result, the CPTs for 'sprinkler' and 'rain' contain only two entries each, in contrast to the four required for 'lawn'.

This translates into drastic savings as the number of variables increases (i.e., n binary variables give rise to 2^n possible combinations of states) and the introduction of BNs fuelled a dramatic rise in the application of the Bayesian framework across a wide range of disciplines and practical applications (Pourret et al., 2008). It is worth noting, however, that, even with BNs, belief revision has a computational complexity (i.e., NP hard, Cooper, 1990) that means that exact calculation rapidly becomes practically impossible as the number of variables increases and approximation algorithms must be used (see e.g., Korb & Nicholson, 2010, for an introduction).

In short, BNs allow the application of Bayes' rule for belief revision across a causally structured Directed Acyclic Graph (DAG) that exploit independence relations in order to allow the calculation of optimal inferences (Pearl, 2009) despite the combinatorial explosion of multi-variate problems. It is thus no accident that in the domain of evidential reasoning specifically, BNs have been successfully used to clarify reasoning errors and fallacies, including evaluating testimony (Madsen et al., 2018, 2019, 2020; Phillips et al., 2018; Pilditch et al., 2020), zero-sum reasoning (Pilditch et al., 2019a), and various errors in legal reasoning (see e.g., Lagnado et al., 2013; Dawid & Evett, 1997; Schum, 1994).

2.1 The limits of intuition

Specifically, the Bayesian framework has been helpful in identifying errors in qualitative and semi-quantitative evaluation rules proposed in the literature. As an example, we take the way argumentation scholars have sought to define putative evaluation rules for argument to cases of testimony (Walton, 2007). In the informal argument literature, there is a long-standing distinction between convergent and linked arguments. The former involves cases where multiple, independent arguments are used to support the same claim. The latter involves cases where arguments somehow 'work together' to generate support. The so-called MAX–MIN rule recommends that for convergent arguments, the strength of the conclusion should equal that of the strongest premise (MAX), whereas for linked arguments, conclusion strength should be determined by the strength of the 'weakest link' (following Pollock, 1995). The MAX–MIN rule has explicitly been endorsed for independent and dependent cases of testimony, but it is normatively appropriate for neither. In both cases, its recommendations conflict with Bayesian prescriptions (see Hahn et al., 2013). The MAX rule ignores any further evidence that is (arbitrarily) less strong than the strongest piece of evidence, so will treat as equivalent the case of a single witness and that single witness plus thousands of additional, independent, only marginally less reliable witnesses. The MIN rule ignores the fact that Bayesian inference across chains (e.g., a three variable BN: A → B → C) can give rise to support for C that exceeds the strength of the weakest link (see Hahn & Oaksford, 2007) and ignores that fact that inference across such a chain may lead to numerically different results depending on the order in which the relative strengths occur (Phillips & Hahn, subm.). The rule consequently does not correctly deal with these kinds of dependent and independent evidence (though related errors can be found in lay reasoners judgments, for convergence/independent evidence see Phillips et al., 2018; 2023; Phillips & Hahn, subm.).

2.2 Variety of evidence

The Bayesian framework has also clarified the variety of evidence thesis. As noted earlier, the variety of evidence thesis has seen widespread discussion within the philosophy of science (Borm et al., 2009; Earman, 1992; Horwich, 1982) and has implications for everything from choice of measurement devices to experimental replication (Oberauer & Lewandowsky, 2019). One insight to come out of the wealth of Bayesian analysis of the variety of evidence thesis (e.g., Bovens & Hartmann, 2003; Fitelson, 1996; Landes, 2021) is that the widespread intuition that more diverse evidence provides stronger support has boundary conditions, and is not always true. As Bovens and Hartmann (2003) demonstrate, whether multiple tests with the same measurement device or whether the same number of tests with multiple devices is preferable depends on the reliability of the device. This has non-obvious implications where the reliability of the device is not fully known. In this case, there are two opposing forces at work: positive tests from two different measuring devices are, all other things equal, more confirming; but repeated positive tests from the same measurement device increase confidence in its reliability (i.e., confidence that it doesn't just generate noise), and that benefits the confirmation a positive test confers on the hypothesis. As a consequence, whether repeated tests with same device or tests with the same outcome taken from different devices increase the probability that the hypothesis at issue is correct depends on the perceived reliability of the device.

Not only do these analytic results refine our understanding of the variety of evidence thesis, they also make clear a fundamental point about independence/dependence. On a superficial analysis, one might not take the above result (or the variety of evidence thesis) to be about dependence at all, because there is a vantage point from which repeated measurements with same device (or replications of the same experiment) might be construed simply as independent events, in the same way as repeated tosses of a coin or throws of a die. That, however, ignores the fact that the outcomes of multiple tests of the same device or the same experimental procedure are likely to exhibit greater correlation because of the shared design (which is the driving factor underlying the variety of evidence hypothesis in the first place). That shared design may be a source of dependence, much in the same way that an actual physical coin may not give rise to wholly independent trials (Bartos et al., 2023) in the same way that an idealised, hypothetical coin does.

This, in turn, highlights two things: The first is that the Bayesian framework (and with it BNs) merely provides a formalism; the extent to which its application affords actual normative insight depends on the adequacy of *the model of the situation* that we build with that formalism. It doesn't magically follow just from the use of the formalism itself. Second, seeking to build such models leads to the identification of conceptually distinct forms of influence or dependence.

2.3 Beyond correlation

Consider awaiting a diagnosis at the hospital regarding whether you have a disease, having had a series of tests conducted independently by two doctors. As in our earlier example of Alex and Jessie, we are put in the position of considering several reports (diagnoses) in relation to a hypothesis (our potential disease). Of critical interest is the relationship between these two reports: are they truly derived independently? Should we account for the two doctors having similar training (shared reliability/backgrounds)? Or using the same test results (shared evidence)? Or have the two doctors passed information between them (direct dependence)?

Bayesian analysis of the kind conducted by Bovens and Hartmann (2003) on the variety of evidence thesis forces one to be clear about these distinctions, because they need to be modelled in different ways. BNs force one to be explicit about different forms of dependence in ways that go beyond purely correlation-based measures (e.g., Berg, 1993) of dependence. This includes counter-factual possibilities and greatly increases the analytic precision over that of a merely correlational framework. Specifically, we can disentangle how different causal structures could explain the same set of observations in one instance, but predict different sets of observations in another. For example, if we consider three sources/reporters (Rep_1, Rep_2, and Rep_3) are all *observed* reporting the same thing, we could take the high correlation between these reports as an indicator of dependence among the sources. However, as shown in Fig. 2 below, this could in fact be due to one or more structural forms.

Although (a) shared evidence, (b) directly conferring (communicating), and (c) shared backgrounds could all give rise to correlations, they may give rise to very different *effects* across circumstances. Lay reasoners fail to appreciate the redundancy entailed by sources sharing evidence (Pilditch et al., 2019b), in line with findings from other frameworks (e.g., Gilliland & Schmitt, 1993; Heller et al., 1992; Kahneman & Tversky, 1973; Maines, 1990), while seemingly more accurate in their estimations of the impact of shared backgrounds (Madsen et al., 2018b, 2019, 2020). Critically, these different forms of dependence can be masked by the same correlation-based pattern of observations.

Furthermore, even direct conferring (b. in Fig. 2) may give rise to difficult to predict patterns of relative advantage for independent vs. dependent cases depending

Fig. 2 Graphical network structure of a reasoning problem. *Note*: Problem structure shows three sources ($Rep_{1–3}$), items of evidence used in their reports ($E_{1–2}$) concerning a hypothesis (H), along with reliabilities for each source ($Rel_{1–3}$), and a shared background (SR) informing those reliabilities

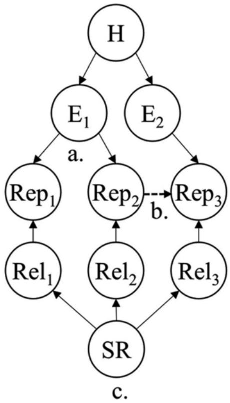

on whether the sources agree in their reports, disagree, or, *only* the receiving source has reported so far. In particular, lay reasoners fail to appreciate that dependencies may normatively give rise to an information *advantage* over independence (Pilditch et al., 2020).

2.4 A closer look at modelling consideration

In order to understand fully the limitations of the Bayesian framework in capturing dependence that will become the focus of the following main part of the paper, it is also important to be fully clear on what the numbers in BNs do and, more importantly, do not represent. Returning to Fig. 1 above, what Bayesian computation allows us to do, in essence, is to calculate, given a bunch of numbers that we specified, what other numbers follow. This is of use because the 'other numbers that follow' may cover a broad range of circumstances, and those numbers themselves may be surprising. In the sprinkler example of Fig. 1 specifically, we needed to enter the respective probabilities associated with the possibility of the sprinkler being on or off, and the probability of rain. These probabilities do not depend on other variables ('sprinkler' and 'rain' do not have any arrows coming in to them), so these are not conditional probabilities but priors. We also have to specify the conditional probabilities of a wet lawn, given the various possible combinations of sprinkler status and rain. Given these specifications in the CPT, the BN then allows us to calculate how probabilities change under a range of observations. Some of these are wholly unsurprising: the network will tell us that, given that the sprinkler is off and there has been no rain, there is a 0.01 probability that the lawn is wet. This is wholly unremarkable, because we basically put that there with our entry in the top, far right, cell of the 'sprinkler' CPT. However, the BN also tells us about the respective probabilities of 'sprinkler' and 'rain' given that the lawn is wet, and these are not probabilities that we specified. Furthermore, the BN tells us how our belief that the sprinkler was on should change when we *additionally* discover that it rained. This additional observation should reduce our degree of belief in the sprinkler having been on, a phenomenon known as 'explaining away', and the details of that phenomenon are sufficiently complex that there exists a sizeable empirical literature examining the extent to which lay reasoners correctly understand such inferences (e.g., Rottman & Hastie, 2016; Tešić et al., 2020). The BN clarifies not just that additionally learning that it rained should influence what our beliefs are, it tells us exactly what that new belief should be, even though we did not specify any direct relationship (conditional probabilities) between ‘sprinkler’ and ‘rain’ (because no direct relationship exists).

How our beliefs concerning the sprinkler should change in those circumstances will depend on the exact numeric values entered in the CPT. Those conditional probability values, like the causal structure itself, reflect our assumptions about the situation. Those assumptions might or might not reflect objective probabilities (to the extent that such probabilities even exist, e.g., Nau, 2001). That is, they might or might not accurately reflect how reliably turning the sprinkler on makes the lawn wet. Most importantly, for what follows, they represent something about the way the sprinkler *is*, not about how the sprinkler *ought to be*. This may seem wholly bizarre to emphasise with respect to a physical device like a sprinkler, but it becomes relevant (as we will see) when we move to considering testimony, because human agents (unlike sprinklers) might take actions that affect their reliability, and this includes (potentially recursively) factoring in informational dependence when forming their beliefs.

We next clarify how and why this sets hard limits on the extent to which the Bayesian framework can capture certain common, naturally arising, cases of dependence, even under extreme idealisation.

3 Bayesian failure: the fundamental problem of social networks

So far, we have considered cases with few evidence sources (measurement devices or human testimony). And we have, so far, seen nothing that would make one suspect this type of formal treatment does not scale. In particular, it might seem possible (at least in principle) to apply it to entire social networks of communicating agents. Both social networks and BNs are graphs, so it seems plausible that we could capture entire social networks with BNs. This seems all the more important given that recent years have made clear that humans, as cognitive agents, are typically agents within a wider social network. Online social networks have arguably changed the scale and density of such embedding, but, in first instance, they have simply made clear the extent to which such embedding is ubiquitous. This means also that dealing with evidential dependence is one of the fundamental challenges human beings need to navigate in everyday life, whether they are actively aware of that challenge or not.

In fact, there are already known limitations on combining the Bayesian framework and social networks. The first, and most obvious, is that in the real-world, we typically do not know the wider structure of our social networks. As an agent in a network, one is trying to form normatively accurate beliefs based on the direct communication one receives from other agents. That communication includes informational dependence that stems from the fact that one’s sources also indirectly communicate *with each other* via paths through the wider social network that link them. In the real world, however, (all) those other agents beyond one’s own sources are unknown, and even less does one know exactly who among them does and does not communicate directly and what they say. So, an appropriate graph cannot be constructed (see also Merdes et al., 2021).

A second practical limitation stems from the computational complexity of inference on Bayesian Networks, already mentioned earlier, in the face of the fact that real-world social networks (in the age of online social media) may comprise billions of nodes (Silberling, 2023). Computational complexity defeats attempts to factor out dependence even under extremely idealised circumstances of a network of Bayesian agents where the entire network structure, agents have uniform priors, and each agent receives a single, initial, piece of private evidence. In this special case, a recursive process of iterative elimination allows agents to gradually infer the body of initial private evidence, but this problem is NP hard (Rahimian et al., 2017). At the same time, the fact that real-world networks have billions of nodes also magnifies the scale of the dependence problem itself, because billions of nodes make for *a lot of potential dependence*.

A third limitation stems from the fact that BNs in the form used here are based on directed acyclic graphs, that is, they must not allow cycles—that is, paths that return to the originating node. Such cycles are already present with bi-directional communication links in a network. This in itself presently rules out the use of BNs to represent wider social networks without further adjustment. There may, however, be technical solutions to the problem of 'loops' or cycles in Bayesian networks (see e.g., Ghahramani, 1997). We do not further consider their adequacy here as there are more fundamental limitations, as we will see.

These three known limitations mean, for example, that Bayesian agent-based models (Madsen et al., 2018a; Olsson, 2011) that simulate communication across social networks make the simplifying (strictly false) assumption that all sources from which a node receives communication are independent. This, demonstrably, can lead to 'double counting' of evidence and resultant beliefs that are more extreme than justified by the actual data in ways that are reminiscent of human errors (see Hahn, 2023 for examples).

The key point of the current paper is that there exist even more fundamental limitations that stem from the fact that, in a sense, BNs are *the wrong kinds of representational objects* for capturing social networks. In short, there is a deeper problem that prevents the mapping, between social networks and BNs, even in highly idealised circumstances where none of the just mentioned limitations plays a role. In other words, the Bayesian framework fails to provide a normative answer for certain simple dependence problems even where all relevant nodes and paths are known, their number is so small as to otherwise allow exact computation, and there are no loops or cycles. In the following, we demonstrate exactly how and why. Specifically, we show how the fundamental features of BNs as representational tools are inadequate for the task at hand. The semantics of BNs and the information they do (and more importantly do not) represent give rise to a number of interlocking limitations.

Returning to our original example let us imagine that instead of receiving information from Alex and Jessie about a potential storm, we instead hear from a third person, Danny. In this version, it is Danny who has spoken to Alex and Jessie, whilst we only hear Danny's report. Crucially, whether Alex and Jessie derived their reports (a) independently of one another, or (b) through a form of dependence (e.g., sources conferring with one another) still intuitively matters just as much as it did before. However, whereas previously we integrated Alex and Jessie's reports ourselves, it is now Danny that does this for us. So, our question now becomes: Given that we have only seen a report from Danny (i.e., Danny has told us there will be a storm tomorrow), how *should we* differentiate our belief in there actually being a storm tomorrow, based on whether Danny got his information through case (a) or case (b)? How should our predictions differ regarding for cases (a) and (b) given Danny's report?

The problems that arise here reflect the just discussed difference between probabilities a BN allows us to calculate, and the probabilities we ourselves need to supply. As we will see, adding in Danny transforms the problem from one in which the BN calculates for us the effects of dependence to one where we have to specify those effects ourselves: the very thing, in fact, that we wanted to use the model for in the first place. We are left stranded because the probabilities that determine the (normative) effect of Danny's testimony in both cases depend on the values entered in Danny's CPT. But those values represent a statement about how Danny actually is, not about what Danny ought to do—in the same way that the CPT for 'wet lawn' implements (assumptions/beliefs about) the way our sprinkler works, not about how it ideally should work in watering the lawn.

To make these points and their implications clearer, we next discuss them with a concrete, reference BN.

3.1 The representation problem

To do so, let us flesh out the two cases: one in which Source 1 (Alex) and Source 2 (Jessie) each receive independent evidence from the world, and the other where they each observed the same piece of evidence (E; e.g., the same weather report).

We can represent this situation with the two BNs in Fig. 3. In each case, there is an underlying state of affairs (Hypothesis = there will be a storm tomorrow) and evidence for that state of affairs dispensed 'from the world". This is not observed by you directly—rather, you receive reports from two sources. Let us assume that the evidence observed (i.e., the weather report(s)) is highly diagnostic, but not perfectly so (P(E|Hypothesis) = 80%, P(¬E|¬Hypothesis) = 80%). Using our main example, the weather reports (E) correctly predict a storm occurring the following day (Hypothesis) 80% of the time. Assume now in the first case (Fig. 4a), there is just one such piece of evidence on which both sources (Alex and Jessie) report (reduplication condition), whereas in the second case each source passes on their own independent observation. Remember also that your beliefs in the hypothesis are represented by the values of the variable Hypothesis, which you are looking down on as an external observer. Your belief in there being a storm tomorrow is represented in the network via the Hypothesis nodes in Figs. 3a and b, but you yourself are *not* a node in the network.

Now let us contrast your (normatively derived) belief in the hypothesis in both cases. In each case, the first report (Source 1/Alex) boosts the posterior degree of belief in the hypothesis from 50 to 68% (Figs. 4a to c for dependent case, and Figs. 3b to d for independent case). In other words, when we have only seen Alex's report, we retain the same degree in belief in their being a storm tomorrow in either case.

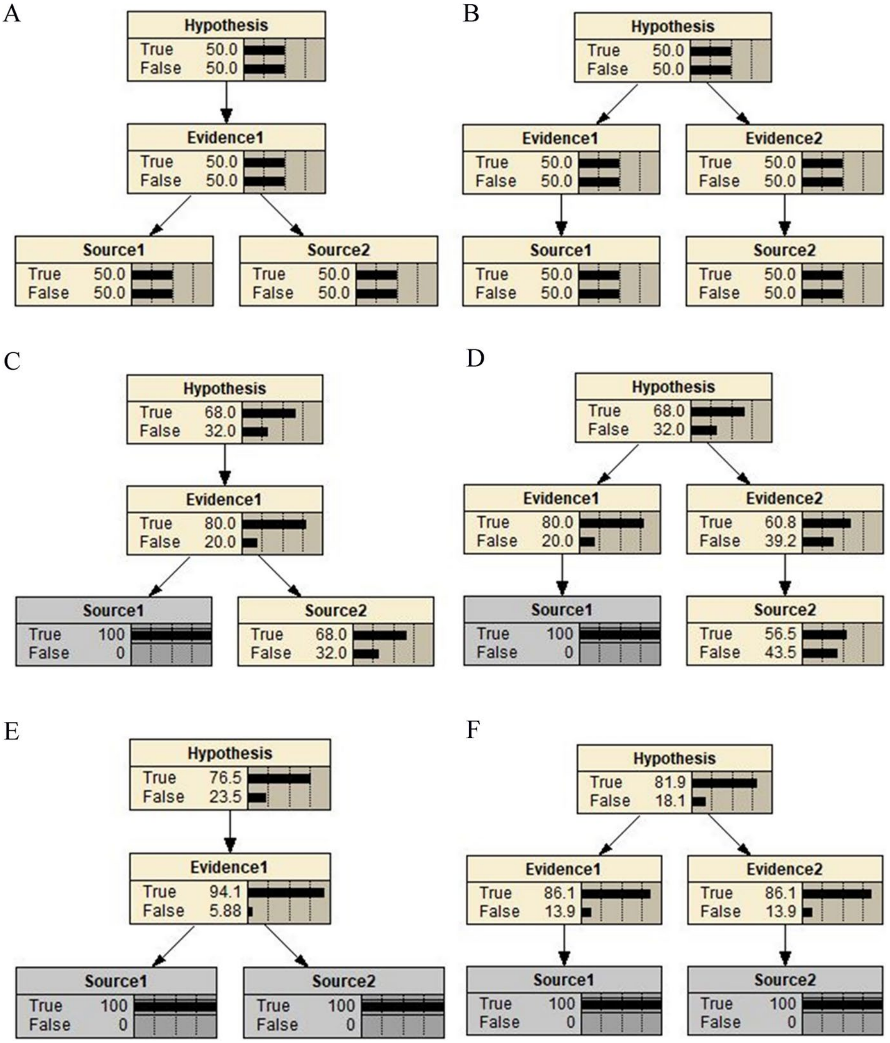

Fig. 3 **A** (left). Shared route of information to sources (no reports observed). **B** (right). Separate routes of information to sources (no reports observed). **C** (left). Shared route of information to sources (first report observed). **D** (right). Separate routes of information to sources (second report observed). **E** (left). Shared route of information to sources (both reports observed). **F** (right). Separate routes of information to sources (both reports observed)

The second report (Source 2/Jessie) then raises it further. But in the independent case, it raises it to 82% (Fig. 3f) whereas in the dependent case it raises it to only 76% (Fig. 3e). In this dependent case there is actually greater confidence in the state of the evidence itself (i.e., given Alex and Jessie are using the same weather report, you can be even more confident in the contents of that report), but this is outweighed in the independent case by the fact that there are two pieces of evidence in agreement (i.e., Alex and Jessie, using different weather reports, have still arrived at the same conclusion).

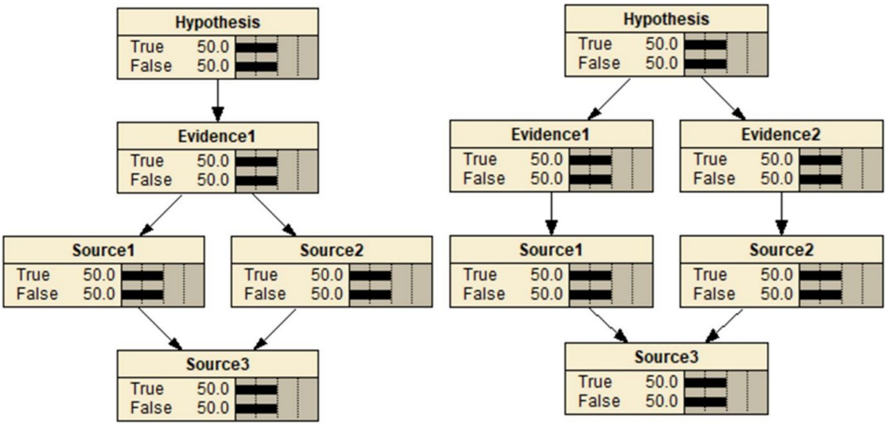

Fig. 4 Shared (left) versus Independent (right) route models, no reports observed

Table 1 Node conditional probability table for source 3

Source 3	S_1 = True, S_2 = True	S_1 = True, S_2 = False	S_1 = False, S_2 = True	S_1 = False, S_2 = False
True	70	50	50	30
False	30	50	50	70

However, what if we now want to introduce a third source (Danny, from the above example)? Consider that both Sources 1 and 2 (Alex and Jessie) report to a third source (Danny), and it is only the report of this third source that is directly observed. First, we need to assign conditional probabilities to the third source as a function of the two reports (from S_1/Alex, and S_2/Jessie). Let us keep this the same in both cases: Source 3 (Danny) has a 70% chance of saying "true" when both the sources it hears from (Alex and Jessie) say "true", and likewise a 70% chance of saying "false" when both say "false". In the remaining cases (i.e., when the two sources disagree), it responds randomly (50/50), as illustrated in Table 1. All other aspects of the networks remain unchanged.

Now, as in our previous example, we then represent this new (network) structure as follows:

If this is what we do, however, we will get *the same degree of confirmation* for the hypothesis in both cases on observing Danny's (Source3's) positive report. This is trivially the case precisely because we *specified the same CPT* for Source3 in both cases. We can obtain a difference that reflects the upstream dependence/independence *only* through those 8 values in Source3's CPT. And, fundamentally, those numbers are input to the Bayesian Network whose determination is wholly outside the BN itself.

Danny/Source3's CPT does not 'see' the rest of the network (and that was in fact the point of capitalising on independence structure to enable BN calculations in the first place, see Sect. 3 above). Danny's report only depends upon the reports of Alex and Jessie—*not the preceding dependent/independent structure that feeds into their reports*. That leaves the CPT blind to the potential effects of a dependency—as Danny is blind to it himself—in any way other than by entering appropriate, dependency reflecting, conditional probabilities for Danny that differ across the two cases. This requires taking an external, "gods-eye", view of the situation (that implicitly incorporates in Danny's CPT structural information beyond his available knowledge). It also, obviously, requires us to know what probabilities *would* appropriately reflect the dependence, which was the very question we were trying to use the BN to address. Simply adding Danny/Source3 transforms the problem from one where the BN tells us something about the normative impact of dependence, to one where a normatively appropriate impact will only emerge because we directly put it in.

Even though we have now realised that the BN itself now no longer helps us derive an appropriate answer to what weight we should give to dependence, we might still hold out hope that there could be a clever solution that would allow us to derive the CPT with which we must equip Danny/Source3 recursively. In effect, can we not capitalise on what we learned in Fig. 3 to determine Danny's CPT in Fig. 4?

In fact, we cannot. We are stopped by a fundamental representation roadblock, which, in fact, is a more deep-rooted problem which the inclusion of Danny has exposed. Put simply, communication across a *social* network (Alex and Jessie to Danny) is not the same (representationally) as probabilistic dependence (evidence to Alex/Jessie)—even though what is being communicated across a social network is evidence. In an agent-based model of communication across social networks (such as Acemoglu et al., 2013; Madsen et al., 2018a; Olsson, 2011), the Bayesian machinery is used to calculate *agents'* degrees of belief in the claim at issue, not whether or not they choose to communicate or what they communicate if they do. Though an agent's degree of belief may influence what they choose to say, there still needs to be a communication rule that determines what agent say and that is not part of their degree of belief per se (see Assaad et al., 2023 for discussion). For example, in the Olsson (2011) model there is a fixed probability that, at a given time step an agent will communicate with a neighbour, and the content of that communication is determined by the agent's 'threshold of assertion'. Specifically, the agent asserts that the hypothesis under discussion across the simulated network is true if the agent's current degree of belief exceeds a threshold probability *p*. Conversely, if the agent's belief is below 1-*p* the agent asserts that the hypothesis is false; and in all other cases, the agent remains silent.

When we think about Bayesian agents in a social network, we are using the Bayesian framework to model beliefs. But in the BN of Fig. 4 neither the node 'Source3' nor its associated CPT represents *Source3's* beliefs. Instead, everything about Source3 represents the beliefs of an *external observer*. The CPT represents *our* assumptions about the evidential value of Source3's report. That might depend on choices Source3 makes (i.e., if Source3 is human, how it aggregates evidence), but that underlying process is nowhere represented. All that is recorded is our assumptions about how reports will be true or false as a function of the states of other variables (Source1 and Source2). Likewise, the marginal probability associated with Source3 at any given time represents *our uncertainty* about Source3, not anything about Source3 itself.

If **we** learn that Alex (Source1) provided a positive report, that will change the probability of the node/variable `Danny' (provided we have not yet received Danny's report), but the probability represents *our* remaining uncertainty about whether or not

Danny will provide a positive report, not anything about Danny's own beliefs or how they might change. Our beliefs about Danny change as we learn the values of other variables both upstream *and downstream* from Danny.

In short, when we are trying to use a BN to form normatively appropriate beliefs about the situation, we are necessarily external to that BN, and that BN represents how our beliefs change in response to different combinations of evidence for us as external observers. Trying to solve the problem of how to respond to Danny's reports recursively would entail having Danny simultaneously be an observer and a node in the network.

And even if we *could* condense Danny's belief formation and communication rules down into the required 8 cell CPT of Table 4, all we would have done is replace actual Danny, with a hypothetical Danny as we think he ought to be. This means that for a wider network, we have simply constructed a counterfactual, hypothetical communication network involving what reports an agent should have received, as opposed to the communication across the network that actually unfolded as it did.

So where does this leave us? Seen from one vantage point, this is simply a Bayesian Network representation problem. It is somewhat bad news for the Bayesian aficionado, but for anyone who is not, this might seem minor. Such a response overlooks that the Bayesian framework is our current *best* tool for elucidating dependence. Moreover, application of that tool revealed the normative implications of dependence (as far as we can track them using the framework) to be oftentimes counter-intuitive and surprising. This makes it bad news all round that the our best tool has such fundamental limitations.

This is made worse by the fact that we need this normative tool (Bayesian or other) to solve not some niche problem, but rather to address what is a fundamental feature of our daily human existence. The communication of testimony is ubiquitous, not just in formal social networks, but in almost all forms of (social) learning (i.e., rarely do we rely on first-hand observation alone). The problems we have identified here in fact mean we cannot understand how impactful dependence (and therefore redundance) is on the strength of such testimony. Therefore, given the omni-presence of various dependence structures (i.e., extremely rarely can we assume complete independence), we simply *cannot, in general, know the strength of testimony accurately*.

The above example is not unusual, nor the consequence of modern phenomena (e.g., social networks), but is in fact present in all forms of communication. Of course, when we consider the implications of this problem, social networks are an increasingly relevant example of the potential impact of dependencies in belief updating. Whether as a tool for the spread of misinformation (e.g., Wang et al., 2019), or as a cascading force of rapid opinion change (e.g., Watts, 2002), the multiply-connected nature of social networks speaks directly to the dangers of underappreciating the impact of dependencies. For example, to underappreciate the degree to which seemingly independent, yet coherent, opinions may be the product of information flowing through the same shared route can lead to overconfidence. This in turn may explain the efficacy of various dis/misinformation phenomena.

4 Conclusions

The question of how we should deal with dependencies is hard to avoid, given an increasingly interconnected world it is also more and more pressing, and the dangers of falsely assuming independence have been well covered in both theoretical (Hogarth, 1989; Pearl, 2009; Rehder, 2014; Schum, 1994) and applied domains (Bex & Prakken, 2004; Kononenko, 1993; Spellman, 2011).

We showed how the Bayesian framework can be used to ward of those dangers by sharpening understanding of when, where and why dependencies matter. Specifically, the framework can be used to identify cases that violate our unrefined, pre-analytic, intuitions by revealing both cases where dependency hurts and cases where dependent measurement devices or testimony actually provide stronger evidence.

With respect to testimony, however, our analysis revealed a fundamental difference between measurement devices and people. People, unlike measurement devices, are typically part of wider social networks of communication in which they actively make choices both about what to believe and what to say. Given the ubiquity of testimony and this basic configuration in our daily lives it thus matters that we identified a surprising limitation to normative insight that arises simply from hearing from Danny instead of his two sources, Alex and Jessie. Specifically, we saw that what seems like an obvious step, using a BN to represent agents in a social network, fails for fundamental representational reasons. BNs represent *our* beliefs about sources, and how those change when we learn that something has been communicated to them. They do *not* represent sources' own beliefs as they are shaped by *communications*. Our representation of such sources (e.g., Danny) reflects *our uncertainty* about them. They do not represent those agents' (e.g., Danny's) own beliefs as we would need to represent in a Bayesian agent-based model of social network communication. The inferential problem we face when confronted with Danny's report is not about what Danny believes, but about the evidential value of Danny's report, and that is a function of what Danny does, not of what he ought to do. To the extent that 'oughts' should factor in to those actions, BNs are no help in making the appropriate determination. Instead, we need that determination to make use of a BN. Thus, although we know beliefs should update in light of testimony, and the evidential value of that testimony will be affected by informational dependencies, capturing that impact as communications pass between sources remains out of reach.

This rules out fine-grained accounts of what a given individual should believe in a dynamically unfolding communication context. It may be possible to circumvent (to some extent) those difficulties by seeking to derive a judgment based on an aggregate, collective output of such sources. The `wisdom of crowds' (Surowiecki, 2005), however, is also fundamentally affected by dependence (e.g., Hogarth, 1978, 1989; Hahn et al., 2018; Jönsson et al. 2015; Becker et al., 2017; Lorenz et al., 2011; Yousif et al., 2019). There has been some success in deriving normative solutions in light of dependence (e.g., Dietrich & Spiekermann, 2013; Ladha, 1993, 1995), but these analyses currently seem far removed from concrete applicability in practical contexts.

Given the centrality of testimony to our lives, and the challenge posed by seemingly ever-increasing inter-connectedness, it is important, in conclusion, to consider possibilities for bypassing a (currently non-existent) normative solution and instead move straight to practical solutions. In particular, what exactly we communicate greatly influences the impact of dependence in practice. Providing context and detail about a source may make transparent when the same information reaches us via different routes and stop us from double counting: being told that John saw a bear in the neighbourhood on Tuesday at 10pm, allows us to avoid over-weighting the evidence when we hear that same thing from Sue, Janice, and Steve, in ways we cannot when they simply tell us that 'someone saw a bear' (see Hahn, 2023). At the same time, societies have at least some control over the density and scale of online social networks. A project such as Facebook, didn't merely move extant networks on line, it sought explicitly to forge novel connections on an unprecedented scale. Potential downsides of informational dependence are something societies might wish to consider with respect to such projects. There may also be other tools that prove useful in trying to minimise the role of dependence. There is some indication, for example, that prediction markets (Ahlstrom-Vij, 2016) involve incentives that make them at least somewhat robust to informational dependence. For example, if the indicator is average market price, and an individual believes themselves to be in possession of highly diagnostic information that suggests it may be inaccurate, that individual should be willing to place a substantial stake on the indicator being wrong. The direction (high/low) and weighting (i.e., size of bet) allows for a signalling function/mechanism that helps compensate for the dependence, and thus one can build a pragmatic solution without requiring a normative one. In short, we may find a range of practical tools or measures to de-correlate individual judgments in order to reduce the deleterious impact of dependence.

Finally, it is entirely possible that the normative tools of the Bayesian framework might yet be expanded in ways that broaden their applicability. In an increasingly interconnected world, where it is progressively more common place that information travels via multiple routes, failing to appreciate the influence of dependencies carries greater risk than ever. Sowe hope that clarifying the challenges may prompt further work to rectify this situation.

Authors contribution All authors contributed to the conceptualisation and writing of this manuscript, and thus the order of authorship should be considered arbitrary. A preprint version of this article has been made available at: https://doi.org/10.31234/osf.io/d2b6e.

Funding UH was supported by a Mercator Fellowship through Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project number 674439.

Declarations

Conflict of interest The authors declare no conflicts of interest.

References

Acemoglu, D., Como, G., Fagnani, F., & Ozdaglar, A. (2013). Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research, 38*(1), 1–27. https://doi.org/10.1287/xxxx.0000.0000

Ahlstrom-Vij, K. (2016). Information markets. *A Companion to Applied Philosophy*, 89–102.

Assaad, L., Fuchs, R., Jalalimanesh, A., Phillips, K., Schoeppl, L., & Hahn, U. (2023). A bayesian agent-based framework for argument exchange across networks. arXiv preprint arXiv:2311.09254.

Bartoš, F., Sarafoglou, A., Godmann, H. R., Sahrani, A., Klein Leunk, D., Gui, P. Y.,... & Wagenmakers, E. J. (2023). Fair coins tend to land on the same side they started: Evidence from 350,757 flips. arXiv e-prints, arXiv-2310.

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences, 114*, 201615978.

Berg, S. (1993). Condorcet's Jury theorem revisited. *European Journal of Political Economy, 9*(3), 437–446.

Bex, F. J., & Prakken, H. (2004). Reinterpreting arguments in dialogue: an application to evidential reasoning. *Legal Knowledge and Information Systems. Jurix*, 119–129.

Borm, G. F., Lemmers, O., Fransen, J., & Donders, R. (2009). The evidence provided by a single trial is less reliable than its statistical analysis suggests. *Journal of Clinical Epidemiology, 62*(7), 711–715.

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

Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review, 124*(3), 301.

Cooper, G. F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence, 42*(2–3), 393–405.

Dawid, A. P. (1979). Conditional independence in statistical theory (with Discussion). *The Journal of the Royal Statistical Society, Series B, 41*, 1–31.

Dawid, A. P. (1980). Conditional Independence for Statistical Operations. *Annals of Statistics, 8*, 598–617.

Dawid, A. P., & Evett, I. W. (1997). Using a graphical method to assist the evaluation of complicated patterns of evidence. *Journal of Forensic Science, 42*(2), 226–231.

Dietrich, F., & Spiekermann, K. (2013). Epistemic democracy with defensible premises 1. *Economics & Philosophy, 29*(1), 87–120.

Earman, J. (1992). *Bayes or Bust?* MIT Press, Cambridge, Mass.

Fenton, N., Neil, M., & Lagnado, D. A. (2013). A general structure for legal arguments about evidence using Bayesian networks. *Cognitive Science, 37*(1), 61–102.

Fitelson, B. (1996). Wayne, Horwich, and evidential diversity. *Philosophy of Science, 63*(4), 652–660. https://doi.org/10.1086/289982

Ghahramani, Z. (1997). Learning dynamic Bayesian networks. In *International School on Neural Networks, Initiated by IIASS and EMFCSC* (pp. 168–197). Springer, Berlin, Heidelberg.

Gilliland, S. W., & Schmitt, N. (1993). Information redundancy and decision behavior: A process tracing investigation. *Organizational Behavior and Human Decision Processes, 54*(2), 157–180.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science, 17*(9), 767–773.

Hahn, U. (2023). Individuals, collectives, and individuals in collectives: The ineliminable role of dependence. *Perspectives on Psychological Science*, 17456916231198479.

Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review, 114*(3), 704.

Hahn, U., Oaksford, M. & Harris, A.J.L. (2013). Testimony and argument: A Bayesian perspective. In Zenker, F. (Ed). *Bayesian argumentation* (pp. 15–38). Springer Library.

Hahn, U., von Sydow, M. & Merdes, C. (2018). How communication can make voters choose less well. In T. T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 1757–1762). Cognitive Science Society.

Hajek, A. (2008). Dutch book arguments. In P. Anand, P. Pattanaik, & C. Puppe (Eds.), *The handbook of rational and social choice* (pp. 173–196). Oxford University Press.

Heller, R. F., Saltzstein, H. D., & Caspe, W. B. (1992). Heuristics in medical and non-medical decision-making. *The Quarterly Journal of Experimental Psychology Section A, 44*(2), 211–235.

Hogarth, R. (1978). A note on aggregating opinions. *Organizational Behavior and Human Performance, 21*, 40–46.

Hogarth, R. M. (1989). On combining diagnostic '"forecasts"': Thoughts and some evidence. *International Journal of Forecasting, 5*, 593–597.

Horwich, P. (1982). *Probability and evidence*. CUP Archive.

Jönsson, M. L., Hahn, U., & Olsson, E. J. (2015). The kind of group you want to belong to: Effects of group structure on group accuracy. *Cognition, 142*, 191–204.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237–251.

Kononenko, I. (1993). Inductive and Bayesian learning in medical diagnosis. *Applied Artificial Intelligence an International Journal, 7*(4), 317–337.

Korb, K. B., & Nicholson, A. E. (2010). *Bayesian artificial intelligence*. CRC Press.

Kruschke, J. K. (2006). Locally Bayesian learning with applications to retrospective revaluation and highlighting. *Psychological Review, 113*, 677–699.

Ladha, K. K. (1993). Condorcet's jury theorem in light of de Finetti's theorem: Majority-rule voting with correlated votes. *Social Choice and Welfare, 10*, 69–85.

Ladha, K. K. (1995). Information pooling through majority-rule voting: Condorcet's jury theorem with correlated votes. *Journal of Economic Behavior and Organization, 26*, 353–372.

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

Landes, J. (2021). The variety of evidence thesis and its independence of degrees of independence. *Synthese, 198*(11), 10611–10641.

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences, 108*(22), 9020–9025.

Madsen, J. K., Bailey, R. M., & Pilditch, T. D. (2018a). Large networks of rational agents form persistent echo chambers. *Scientific Reports, 8*(1), 12391.

Madsen, J. K., Hahn, U., & Pilditch, T. D. (2018b). Partial source dependence and reliability revision: the impact of shared backgrounds. In T. T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 722–727). Cognitive Science Society.

Madsen, J. K., Hahn, U., & Pilditch, T. (2019). Reasoning about dissent: Expert disagreement and shared backgrounds. In *Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 41)*.

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.

Maines, L. A. (1990). The effect of forecast redundancy on judgments of a consensus forecast's expected accuracy. *Journal of Accounting Research, 28*, 29–47.

Merdes, C., Von Sydow, M., & Hahn, U. (2021). Formal models of source reliability. *Synthese, 198*, 5773–5801.

Nau, R. F. (2001). De Finetti was right: Probability does not exist. *Theory and Decision, 51*, 89–124.

Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Prentice Hall.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101*(4), 608.

Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review, 26*, 1596–1618.

Olsson, E. J. (2011). A simulation approach to veritistic social epistemology. *Episteme, 8*(02), 127–143.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Morgan Kaufmann.

Pearl, J. (2009). *Causality. models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.

Phillips, K., Hahn, U., & Pilditch, T. D. (2018) Evaluating testimony from multiple witnesses: Single cue satisficing or integration? In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*.

Phillips, K., Hahn, U., & Pilditch, T. D. (2023). Evaluating testimony from multiple witnesses: Consistent undervaluing and selective devaluing of corroborating reports. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 45, No. 45).
Pilditch, T. D., Fenton, N., & Lagnado, D. (2019a). The zero-sum fallacy in evidence evaluation. *Psychological science, 30(2)*, 250–260.
Pilditch, T. D., Hahn, U., & Lagnado, D. (2019b). Shared Evidence: It all depends…. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 41).
Pilditch, T. D., Hahn, U., Fenton, N., & Lagnado, D. (2020). Dependencies in evidential reports: The case for informational advantages. *Cognition, 204*, 104343. https://doi.org/10.1016/j.cognition.2020.104343
Pollock, J. L. (1995). *Cognitive carpentry: A blueprint for how to build a person*. MIT Press.
Pourret, O., Na, P., & Marcot, B. (Eds.). (2008). *Bayesian networks: A practical guide to applications*. Wiley.
Rahimian, M. A., Jadbabaie, A., & Mossel, E. (2017). Complexity of Bayesian belief exchange over a network. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)* (pp. 2611–2616). IEEE.
Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive Psychology, 72*, 54–107.
Rottman, B. M., & Hastie, R. (2016). Do people reason rationally about causally related events? Markov violations, weak inferences, and failures of explaining away. *Cognitive Psychology, 87*, 88–134.
Schum, D. A. (1994). *The evidential foundations of probabilistic reasoning*. Northwestern University Press.
Silberling, A. (2023) https://techcrunch.com/2023/07/26/facebook-3-billion-users/
Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual Review of Psychology, 66*, 223–247.
Smith, A. E., Ryan, P. B., & Evans, J. S. (1992). The effect of neglecting correlations when propagating uncertainty and estimating the population distribution of risk. *Risk Analysis, 12*(4), 467–474.
Soll, J. B. (1999). Intuitive theories of information: Beliefs about the value of redundancy. *Cognitive Psychology, 38*(2), 317–346.
Spellman, B. A. (2011). Individual reasoning. In B. Fischhoff & C. Chauvin (Eds.), *Intelligence Analysis: Behavioral and Social Scientific Foundations.* National Academies Press.
Spohn, W. (1980). Stochastic independence, causal independence, and shieldability. *Journal of Philosophical Logic, 9*, 73–99.
Surowiecki, J. (2005). *The wisdom of crowds*. Anchor.
Tešić, M., Liefgreen, A., & Lagnado, D. (2020). The propensity interpretation of probability and diagnostic split in explaining away. *Cognitive Psychology, 121*, 101293.
Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic bulletin & review, 14(5)*, 779–804.
Walton, D. (2007). *Witness testimony evidence: Argumentation and the law*. Cambridge University Press.
Walton, D. N. (2004). *Relevance in argumentation*. Erlbaum.
Wang, Y., McKee, M., Torbica, A., & Stuckler, D. (2019). Systematic literature review on the spread of health-related misinformation on social media. *Social Science & Medicine, 240*, 112552.
Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America, 99*(9), 5766–5771.
Yousif, S. R., Aboody, R., & Keil, F. C. (2019). The illusion of consensus: A failure to distinguish between true and false consensus. *Psychological Science, 30*(8), 1195–1204.