

Scientific Theories as Bayesian Nets: Structure and Evidence Sensitivity

Patrick Grim, Frank Siedl, Calum McNamara, Hinton Rago, Isabell Astor, Caroline Diaso, and Peter Rynor*

Quine (1951) painted a picture of science as a web of belief:¹

The totality of our so-called knowledge or beliefs, from the most casual matters of geography and history to the profoundest laws of atomic physics or even of pure mathematics and logic, is a man-made fabric which impinges on experience only along the edges. Or, to change the figure, total science is like a field of force, whose boundary conditions are experience. A conflict with experience at the periphery occasions readjustments in the interior of the field. Truth values have to be redistributed over some of our statements. Reevaluation of some statements entails reevaluation of others, because of their logical interconnections.... (39)

As Brian Skyrms and Karl Lambert (1995) note, “however attractive this picture may be, Quine does not offer any methodology for modeling and mapping the networks of belief...” (139).

Bayesian nets offer precisely such a model, and do so in many of the respects called for by Skyrms and Lambert: “We believe that the best framework for a precise realization of these ideas is the theory of personal probability. The question then arises how to map a network of degrees of belief in a way which reveals the weak and strong resistances to disconfirmation and more generally how need for revision tends to be accommodated by the network.... Based on the conditional probability structure, notions of invariance, independence, and conditional independence all play a role in determining the place of a statement on the web of belief” (139-40).

*We are grateful to Gordon Belot, Josh Hunt, Jim Joyce, Mark Newman, Brian Skyrms, and two anonymous referees for help and encouragement. Patrick also owes a note of thanks to the Simulations of Scientific Inquiry research network, funded the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 426833574.

1. See also Quine and Ullian (1970).

‘Webs of belief’ in general, personal or cultural, scientific or non-, may take various forms, demanding various patterns of connection and updating. We don’t intend the models offered here to represent them all. We concentrate on a particular type of web—that characteristic of much of science, with emphasis on two characteristic relations: inference and confirmation. Bayesian nets, we propose, offer a particularly promising model of inference and confirmation within scientific theories.²

Bayesian approaches to confirmation, of which there are important variations, have been widely applauded as alternatives to their twentieth-century predecessors in the philosophy of science, which were couched largely in classical logic (Skyrms, 1987; Easwaran, 2011; Pettigrew, 2016; Crupi, 2020). But the success of Bayesianism, contrasts with its predecessors, and debates over its variations have largely proceeded using examples of a single piece of evidence, new or old, and its impact on a single hypothesis. Here we emphasize the *networks* of propositions that constitute scientific theories, a theme central, though in relatively vague qualitative form, in a very different tradition in philosophy of science (Kuhn, 1969; Lakatos, 1969; Laudan, 1978) and in an extensive history of conceptual maps scattered across disciplines (Axelrod, 1976; Aguilar, 2013; Kosko, 1986; Papageorgiou and Stylios, 2008; Hobbs et al., 2002; van Vliet et al., 2010; Soler et al., 2012; Cakmak et al., 2013; Jetter and Sperry, 2013; Findlay and Thagard, 2014; Homer-Dixon et al., 2014; Thagard, 2015).

Other efforts have been made to model networks of *agents* updating on information from other agents, with various degrees of reliability or trust (Bovens and Hartmann, 2002; Bovens et al., 2003; Olsson, 2013; Olsson and Vallinder, 2013). This is again a different picture and a different target from the one we pursue here. Our closest predecessors are Henderson et al. (2010) and Climenhaga (2019, forthcoming) The first of these is restricted to strictly linear hierarchical models of levels of abstraction with a focus on Bayesian issues of simplicity. Climenhaga represents explanatory relations between propositions as Bayesian nets in ways that complement our work here, informally in Climenhaga (forthcoming) and more formally, but with a focus on the specific question of which probabilities determine the values of other probabilities, in Climenhaga (2019). A more complete quantitative model of scientific theories, tracing the dynamics of evidence, implication, and confirmation percolating through branching networks of propositions, has not yet been fully drawn.

In this paper we concentrate on the structure of scientific theory, and how that

2. A very different model of Quinean webs of belief as NK landscapes is offered in passing in Alexander et al. (2015).

structure determines sensitivity to evidential impact and credence change. Evidence of a given strength can have far greater impact at one node than at another in the same structure. Webs of belief with different structures can be differentially vulnerable or resistant to the impact of evidence at a given level, or the impact of a pattern of evidence over time. A study of network structure can thus reveal what Skyrms and Lambert hoped for: “the weak and strong resistances to disconfirmation” and more generally “how need for revision tends to be accommodated by the network.”

§2 outlines a Bayesian network model of scientific theories, and we give examples of the differential impact of evidence at different nodes in §3. In §4 those results are generalized using a graphic measure of differential evidence impact at different nodes, with an analysis through examples in §5 of how evidence works: the interplay of network factors in the impact of evidence. To this point, for the sake of simplicity, the theoretical structures considered have been limited to poly-trees. In §6 the treatment is expanded to directed acyclic graphs in general. The search for generalizations regarding influential nodes in networks of various sorts is extensive across disciplines, incomplete, and ongoing (Kempe et al., 2005; Chen et al., 2013; Lawyer, 2015; Bao et al., 2017; Li et al., 2018; Wei et al., 2018; Hafiene et al., 2019; Champion and Elkan, MS). Ours is a particular form of that search, geared to a particular kind of evidential influence in theoretical credence networks in particular.

We conclude in §7 with further emphasis on philosophical implications of model simulations. Basic conclusions, though perhaps qualitatively intuitive, are captured here in a quantitative model open for analysis. Our attempt in the paper as a whole is to take some first steps in understanding the contribution to epistemic sensitivity and significance made by the network structure of our theories.

1 A Bayesian Model of Scientific Theory

Bayesian networks have been developed, interpreted, and widely applied as models of causal relations between events (Pearl, 1988, 2009; Spirtes et al., 2000; Spirtes, 2010). But it is clear that they can equally well be taken as models of causal inference between propositions descriptive of those events: as theories. One form that scientific theories take is precisely this: a propositional representation of a causal network, between either types or tokens of events. We can therefore study structural aspects of at least one type of scientific theory by applying structural lessons from Bayesian nets.

A partial reconstruction of a causal theory for the failure of the 17th Street

canal levee in New Orleans during Hurricane Katrina is shown in Figure 1. A major assumption in many applications of Bayesian nets with an eye to causal inference is that all causal factors, or at least all relevant causal factors, direct or latent, are included in the representation. This is certainly *not* true of the reconstruction here—a partial reconstruction that captures the spirit though not the detail of a full causal theory.

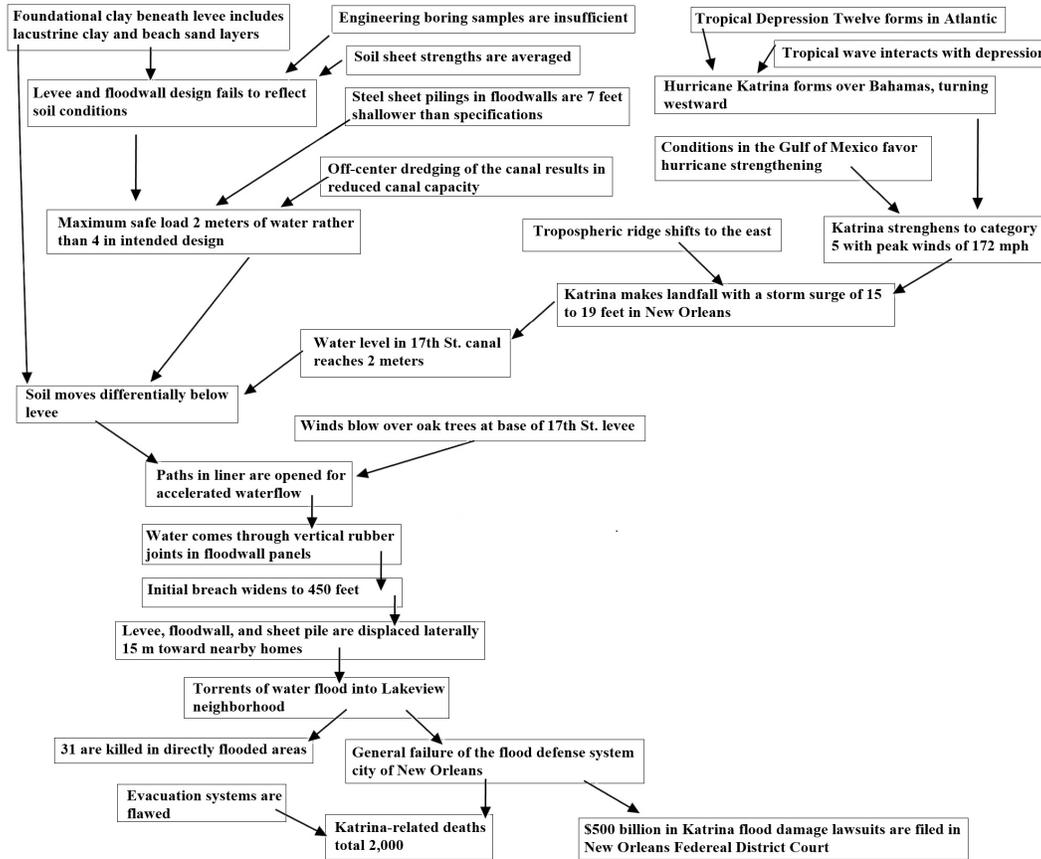


Figure 1: A causal theory for the failure of the 17th Street canal levee in New Orleans during Hurricane Katrina. Sources: Anderson et al., 2007; Panel 2007, Bea, 2008; Rogers et al., 2008; Boyd, 2010.

Lessons from causal models can also be generalized (Schaffer, 2016; Climenhaga, 2019). The propositions within a scientific theory need not be descriptions of events, and the relations between them need not be those of causal inference. From fundamental and more general hypotheses within a theory, which might be envisaged as root nodes, multiple layers of derivative and more specific hypothe-

ses may be inferred (Henderson et al., 2010). That inference may be probabilistic, from general grounding hypotheses to the more specific hypotheses they ground, modellable by conditional probabilities precisely as in Bayesian nets—though here our arrows represent not ‘ x causes y ’ but ‘the credibility of y is grounded in the credibility of x ’ and often ‘ x explains y .’ In the other direction, evidence for more specific or applicational hypotheses can serve to confirm the more general hypotheses from which they can be inferred. This direction is often clearest in terms of disconfirmation: to what extent would x be disconfirmed were its implication y disconfirmed? Climenhaga (forthcoming) outlines a network approach of this type informally, intended to include inference to the best explanation, enumerative induction, and analogical inference.

A partial reconstruction of a theoretical structure of this type regarding the COVID-19 pandemic is shown in Figure 2. As in the causal case, a full Bayesian net representation would require all relevant grounding nodes. This is certainly *not* true of the reconstruction here—a partial reconstruction that captures the spirit though not the detail of a full theoretical structure.

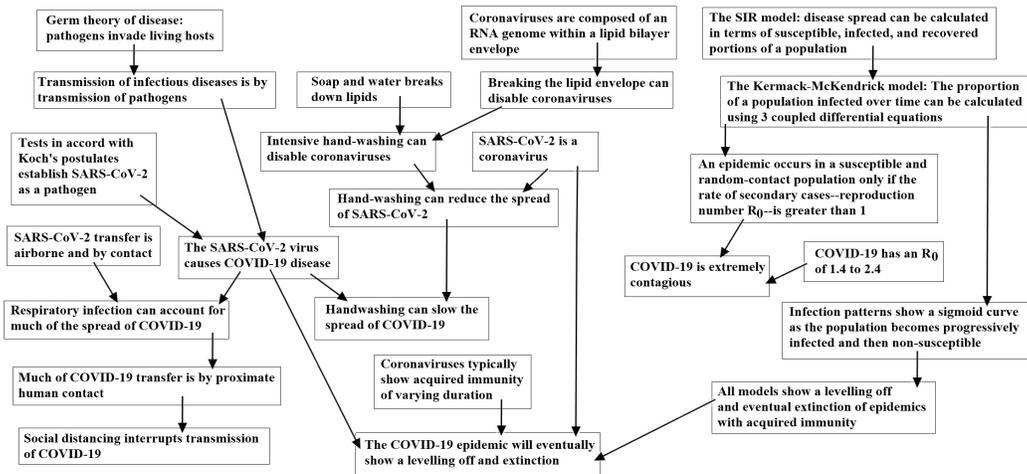


Figure 2: A reconstruction of foundational theory for the COVID-19 pandemic. Sources: Kermack et al., 1927; Anderson and May, 1979; Hassan et al., 2020.

Classical twentieth century models of inference and confirmation were couched largely in terms of classical logic. Bayesian models are seen as an alternative. But there is a general pattern that the two clearly share. In classical models, evidence for a hypothesis is evidence for that which it entails. In the current model, we use inference rather than entailment, cashed out in Bayesian terms of priors and conditional probabilities. In classical models, evidence for a hypothesis offers con-

firmation for hypotheses which entail it. Here, we take evidence for a hypothesis to offer confirmation for hypotheses from which it may be inferred, but confirmation appears in a quantitative Bayesian rather than qualitative Hempelian form. Although in accord with the spirit of classical philosophy of science, an image of scientific theories as Bayesian nets offers a far more quantitative and far more nuanced view of the basic relations of theoretical inference and evidential confirmation.³

We model a scientific theory as a directed acyclic network.⁴ Nodes represent proposition-like elements which carry credences (degrees of belief). It seems natural to speak of nodes as carrying ‘beliefs,’ and we will, though it is no part of our effort to defend that nomenclature literally nor to enter the debate as to what degree of credence or commitment, absolute or contextual, qualifies an item as a full ‘belief.’ Spohn (2012) emphasizes the difficulties posed by the Lottery Paradox in the latter regard. It is sufficient for the abstract purposes of our model that the nodes represent elements of a ‘theoretical system’ or ‘web of belief’ that carry degrees of credence or commitment. Our nodes take values in the open interval $(0,1)$, excluding 0 and 1 themselves, though variations which include these numbers are also possible.

The links of our model represent the connections between the claims within a scientific theory. What is crucial is that the elements in the network are linked in such a way that credence change at one point produces credence change at the points with which it is linked. Our directed links $x \rightarrow y$ carry weights as conditional probabilities, allowing us to update both y conditional on a given credence at x and x on the basis of evidence at y . Our model networks thus accord with the Markov condition: conditional on their parents, nodes are independent of all non-descendants.

Our attempt is to model scientific networks or theoretical webs of belief as Bayesian nets, percolating credence changes at nodes through a given network structure. Some major limitations should be noted, both within this model and beyond.

First, within the model, simplifying assumptions regarding conditional probability and credence assignments are noted throughout. We work with static network structures in the models explored here: neither nodes nor links are added

3. One thing a Bayesian approach does not do, unfortunately, is solve classical paradoxes of confirmation such as Goodman’s grue puzzle (Goodman, 1955). See Bandyopadhyay et al. (2016) for further discussion.
4. Contact of our approach with either ‘syntactic’ or ‘semantic’ models of scientific theories is remote. For the relation to these, and superior features of a Bayesian model, see Skyrms (1984) and Hartmann (2008).

or subtracted. We also work with single snapshots of evidence impact; we leave to further work the promising study of dynamic network change with iterated evidence over time, including rolling changes in both node credences and conditional probabilities on links. In all these regards it must be admitted that what we offer here are merely first steps in modeling scientific theories as Bayesian nets.

Second, beyond this model, it must be admitted that there are probabilistic network alternatives to straight Bayesian nets, well worthy of exploration. Our hope is that the exploration of Bayesian nets in particular may motivate further work on these alternatives as well.

2 Differences in Evidence Impact: Initial Examples

We start with a simple illustration of the relative importance of differently positioned nodes in a theoretical structure. In some cases evidence may affect just one belief or several closely related beliefs in a significant way, decaying quickly in influence across the network. In other cases, evidence of the same strength regarding another element with a scientific theory can cascade with a far stronger influence throughout the network.

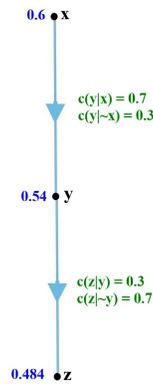


Figure 3: A simple Bayesian network

Figure 3 shows a simple example of a theoretical structure modeled as a Bayesian network, with credences and conditional probabilities (marked as conditional credences) in place. We offer a piece of evidence of strength $[0.75, 0.25]$ at node y with prior credence 0.54, indicating that such a piece of evidence has a likelihood of 0.25 should y be true, a likelihood of 0.75 should y be false. On standard Bayesian conditioning, our credence at y given that piece of evidence

changes from 0.54 to approximately 0.28:

$$\begin{aligned} p(y | e) &= \frac{p(e | y) \cdot p(y)}{(p(e | y)) \cdot (p(y)) + (p(e | \sim y) \cdot p(\sim y))} \\ &= \frac{0.25 \cdot 0.54}{(0.25 \cdot 0.54) + (0.75 \cdot 0.46)} \\ &= 0.28125 \end{aligned}$$

But of course, a change in credence at y demands a change in credences downstream as well. Given a new credence of 0.28 at y with an indicated conditional probability of z given y , and assuming the evidence e affects credence at z only by way of changed credence at y , our revised credence at z given evidence e at y changes from 0.484 to 0.587:

$$p(z | e) = (p(z | y) \cdot p(y | e)) + (p(z | \sim y) \cdot p(\sim y | e))$$

by independence of z and e conditional on y :

$$p(z | e) = (0.3 \cdot 0.28) + (0.7 \cdot 0.72)$$

In an extended network, with the same independence assumption at each step, changes would continue downstream in the same manner.

Credence change percolates upstream as well. Using

$$p(x | y) = \frac{p(y | x) \cdot p(x)}{p(y)}$$

from Bayes, with our initial priors and conditional probabilities,

$$\begin{aligned} p(x | y) &= \frac{0.7 \cdot 0.6}{0.54} \\ &= 0.78. \end{aligned}$$

Using

$$p(x | \sim y) = \frac{p(\sim y | x) \cdot p(x)}{p(\sim y)},$$

our initial values give us

$$\begin{aligned} p(x | \sim y) &= 0.3 \cdot 0.6046 \\ &= 0.39. \end{aligned}$$

Again, making use of an assumption of independence, the updated credence for x given our new value for y on evidence e at y is

$$p(x | e) = p(x | y) \cdot p(y | e) + p(x | \sim y) \cdot p(\sim y | e)$$

(by independence of x and e conditional on y)

$$p(x | e) = (0.78 \cdot 0.28) + (0.39 \cdot 0.72)$$

Thus credence at parent node x changes from 0.6 to 0.5.⁵

On this picture, the direct impact of evidence on a given node ramifies one step downward to its immediate descendants and one step upward to its parents. With a screening-off assumption at each step, revised credence values for the descendants of those descendants—and parents of those parents—can be calculated in the same fashion. What we have in effect is a game of telephone, tracing revised credences percolating throughout the structure of a scientific theory.⁶

We can use the absolute difference between prior and posterior credences at each node as a simple measure of credence impact at each node. Using that measure in this example, a piece of evidence of strength [0.75, 0.25] at node y has an impact of 0.26 at node y and of approximately 0.1 at the other nodes, giving a total network impact of 0.46.

Starting with the same priors, consider now evidence of the same strength at root node x instead. In that case, credence at x changes from 0.6 to 0.33, credence at y from 0.54 to 0.43, and credence at z from 0.48 to 0.526. The total network impact is 0.42. In this network, with these conditional probabilities and priors, evidence impact at the central node y dominates the impact of the same evidence at the root node x .⁷

-
5. Technically, our updating is by Jeffrey conditioning (Jeffrey, 1965; Shafer, 1981) rather than strictly Bayesian, since beyond the initial point of evidence updating is in terms of partitions (here binary) and using credences other than 1. We follow common practice in referring to the Jeffrey generalization as Bayesian.
 6. As indicated below, backtracking will also occur. In a ‘star’ or ‘hub’ tree with one root and many leaf nodes, change at one leaf of the tree will affect credence upward at the root node, which will then in turn affect credence downward at a second leaf.
 7. Here we use a difference measure between posterior and prior credences, $| p(x | e) - p(x) |$, as our measure of impact. This carries over to our use of Brier divergence below. An alternative would be to use a ratio measure $p(x | e)/p(x)$. In both this example and in general that alternative measure would of course change the absolute values of evidence impact, but in neither this case nor in general would it appear to make any difference in the *relative* impact of evidence at different nodes: at which nodes a given strength of evidence has the greatest network impact. The question of appropriate metrics for evidence impact is closely allied with the question of appropriate metrics for degree of confirmation. Crupi et al. (2007) offer a very useful overview of the latter.

A similar story holds if evidence of the same strength is delivered at node z . In that case credence change at node z is from 0.48 to 0.23, the value of y becomes 0.637 from 0.54 and the value of x becomes 0.637 from 0.6 with a total network impact of 0.38.

Given the structure of this simple network, the patterns of belief-to-belief influence modeled by conditional probabilities, and with these initial credences in place, it is the central node y that carries the most network-wide impact for a piece of evidence of the same strength.

Different network structures, with different conditional links and different priors, will clearly exhibit different evidence sensitivities. A second example is offered in Figure 4.

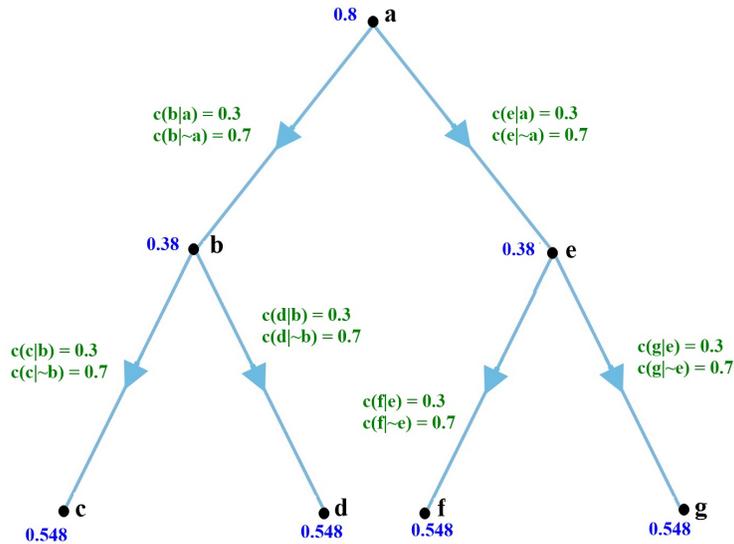


Figure 4: A second simple Bayesian network

Initial credences and conditional probabilities are shown in the graph. With a piece of evidence of strength $[0.9, 0.1]$ at node a —indicating a piece of evidence with a conditional probability of .1 if a is true, 0.9 if a is false—our credences change to 0.3 at the root node, 0.579 at each central branch, and 0.469 at the leaves: a change of .5 at the root node, 0.199 at the central branches and 0.079 at the leaves. The total change for the network is 1.214, or an average change of 0.173 per node. That same strength of evidence at g gives us a total change in credences of 0.602, or an average change of 0.086 per node, less than half the change at the root. Evidence of that strength is clearly of greater impact at the top than the bottom—at the root than at the leaf.

But the situation is different if our evidence is of strength $[0.1, 0.9]$ rather than $[0.9, 0.1]$ instead. In that case the total change at the root node a is 0.418, an average change of merely 0.059 per node. Evidence of that same strength at leaf node g gives us a total change of 0.623—an average change of 0.089. Evidence with the conditional probabilities reversed thus has a roughly 50% greater impact in this network at a leaf node g than at the root node a .

What even these simple examples demonstrate is disparity in the effect of evidence given the structure of a theoretical network, the conditional probabilities of its links, and initial priors. Given the same structure and conditional probabilities, moreover, which is the most evidentially influential node may well depend on the character of the evidence itself: in this example, $[0.1, 0.9]$ rather than $[0.9, 0.1]$. At what point is a scientific theory most vulnerable to change in light of evidence? Using a Bayesian model, we have to conclude that the answer can depend on all these factors: network structure, conditional probabilities, prior credences, and the character of the evidence itself. The role that each has to play, and their interaction, are outlined roughly below.

3 A Graphic Portrait of Evidence Impact

What we have offered above are single-point snapshots of the differential impact of evidence of the same strength at different nodes within a network. We can envisage differential node impact more generally with graphs such as that in Figure 8, illustrating impact in the network shown in Figure 7.

Our graphing conventions differ from those used above. In our examples, evidence strength was indicated in terms of two conditional likelihoods such $[0.8, 0.1]$ at a node h —indicating a piece of evidence with a probability of 0.1 if h is true, or 0.8 if h is false. In fact, this is more information than is needed to determine the impact of evidence on the entire network. The ratio of these likelihoods—in this case $1/8$ —fully determines evidential strength. Hence, it does not matter whether our evidence is $[0.8, 0.1]$, $[0.4, 0.05]$ or $[0.08, 0.01]$; both the immediate effect on the node in question and its over-all impact on the network as a whole will be the same. Therefore, we take the likelihood ratio as a concise description of evidence strength, used on the horizontal axis of our graph in Figure 5. A log scale is used to illustrate the symmetry between what we can think of as positive evidence, with likelihood ratios in $(1, \infty)$, and negative evidence, with likelihood ratios in $(0, 1)$. We choose base 2 somewhat arbitrarily.

We used a simple total of changes across a network and the average change per node in our calculations above. A more elegant measure of network change

is Kullback-Leibler divergence, a relative entropy measure given by:

$$D_{\text{KL}}(p, q) = \sum_{i=1}^n p(A_i) \cdot \log(p(A_i | q(A_i)))$$

where p and q are credence functions (probability measures), and the A_i are the propositions (events) over which these functions are defined. Kullback-Leibler divergence gives the difference between probability measures in terms of a sum of products, namely the product of the probability of the A_i according to p , and the log of the probability of the A_i according to p conditional on the probability of A_i according to q . It can thus be interpreted as the expected logarithmic difference between distributions p and q , with the expectation taken according to p . Kullback-Leibler divergence averaged over nodes is taken as our measure on the y axis in Figure 5.

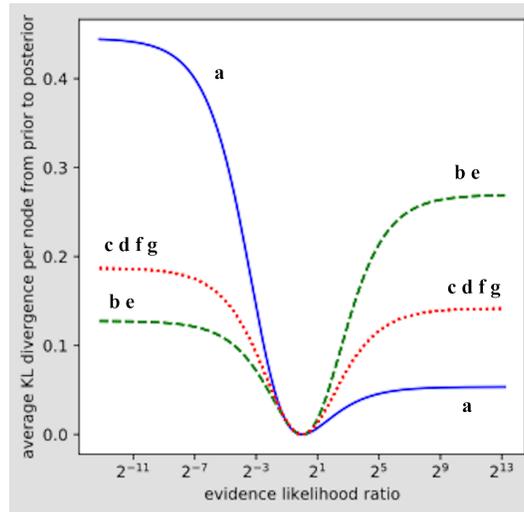


Figure 5: Node-averaged Kullback-Leibler divergence for evidence strength (in terms of exponent of likelihood ratio) at different nodes of the network in Fig. 4.

The graph is read in terms of the Kullback-Leibler divergence across the entire network for evidence of a specific strength at marked nodes. On the right side of the graph, a piece of evidence with strength ratio of 2^3 or $8/1$ at node a , for example, results in a Kullback-Leibler divergence for the network as a whole of approximately 0.03. A piece of evidence of that same strength at nodes c, d, f or g results in a Kullback-Leibler divergence of approximately 0.07 for the network as a whole. The network impact of evidence of that strength at nodes b or e is approximately 0.13. On the left side of the graph, a piece of evidence with

strength 2^{-3} or $1/8$ has a Kullback-Leibler network impact of 0.07 at nodes b and e , 0.09 at nodes c, d, f , or g , and approximately 0.14 at root node a .

An alternative measure for network change, less familiar in some formal disciplines but more common in the philosophical literature, is Brier divergence, proportional to a squared Euclidean distance measure (Joyce, 2009; Pettigrew, 2016). The difference between prior and posterior probability distributions is taken to be the mean of the squared difference between the prior and posterior probability attached to each proposition in the network. Formally:

$$D_{\text{Brier}}(p, q) = \frac{1}{n} \sum_{i=1}^n (p(A_i) - q(A_i))^2$$

Results for Brier divergence, read in the same manner, are shown in Figure 6.

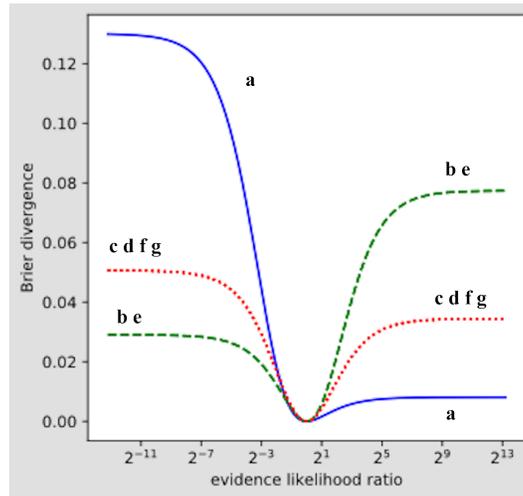


Figure 6: Node-averaged Brier divergence for evidence strength (in terms of exponent of likelihood ratio) at different nodes of the network in Fig. 4.

Kullback-Leibler divergence is a measure of probability distance between the distributions of individual nodes, averaged in Figure 5 to obtain a divergence metric for the network as a whole. Brier divergence is defined over a set of nodes, so it naturally serves as a network divergence metric.⁸

These more complete graphical analyses underscore the conclusions drawn above: that epistemic impact across a network of evidence at a node z can depend on both z 's place in the network and on the character of the evidence, gauged

8. In general, Kullback-Leibler is a multiplicative score, whereas Brier is additive; it should also perhaps be noted that Brier is symmetric, whereas Kullback-Leibler is not.

in terms of its likelihood ratio: the probability given that proposition z is true over its probability given that proposition z is false. In this network, with these priors, the strongest impact for evidence more likely if a node's value is false is at the root node. The strongest impact for evidence more likely if a node's value is true is at the central nodes of the network. Evidence at the leaf tips, within these parameters for evidence of either character, falls in between.

We can draw two clear though simple conclusions on the basis of the work to this point, pursued further in the details of section 5.

The first conclusion is that in this case, as in network studies in general, the impact of evidence at a node depends on the position of that node within the theoretical structure. That is clear 'vertically' from the graphs in Figures 5 and 6, for example: values at any point above or below an exponent of 0 vary with node position.

A second conclusion, in contrast to other network studies, is that the Bayesian structure of these networks is crucially important. This is clear 'horizontally' from the graphs in Figures 5 and 6. Any network measure that fails to incorporate the crucial factor of the character of impact at a node, in terms of likelihood ratio, will fail to identify the most epistemically influential nodes. That impact changes quite dramatically from the left to the right of graphs such as those in Figures 5 and 6, though the network structure, all conditional probabilities, and all priors remain the same.

4 How Evidence Works: Interactive Factors in Scientific Networks

Results like those found in section 4 offer the prospect of being able to gauge epistemically sensitive nodes within the structure of a scientific theory: those elements at which evidence of a particular character would have a particular impact in terms of credence change across the network as a whole, and those points at which a scientific theory would be most evidence-sensitive to change. Establishing precisely how evidence impacts a theoretical network, however, quickly becomes a challenging task.

There has been a disciplinary widespread search, for various purposes—social, technological, and epidemiological—for principles governing most influential nodes within a network (Kempe et al., 2005; Chen et al., 2013; Bao et al., 2017; Li et al., 2018; Wei et al., 2018; Hafiene et al., 2019; Champion and Elkan, MS). Essentially all of this work, however, has been on networks far simpler than the Bayesian models we envisage here: on undirected networks without weighted links, for ex-

ample. To the extent that consensus has emerged on network measures for most influential nodes—a consensus is far from complete—those measures will also prove too simple for Bayesian models.

Vulnerability and robustness of network structures have been studied in terms of the removal of nodes in undirected networks, with applied examples such as circumstantial mechanical failures within electrical grids, evolutionary impact on specific species within an ecosystem, and terrorist attacks on airline or internet hubs (Albert et al., 2000; Newman, 2018). For removal of nodes at random, it is scale-free and preferential attachment networks—those with a small proportion of highly connected ‘hubs’—that prove the most resilient against attack. Random removal is most likely to hit something other than a hub, doing relatively little damage. These are also the most vulnerable networks for targeted attacks, however, such as deliberate acts of informed sabotage: removal of a few carefully chosen hubs can do a great deal of damage. We can certainly expect something like a hub effect in our Bayesian networks. In a downward direction, credence change in nodes with more children can to that extent be expected to have an effect on a wider number of nodes and thus a more pronounced effect on the network as a whole. But here the character of that effect, and its extent, must also take into account the conditional probabilities on links to child nodes—not merely the existence of a link—and the character of evidence at the parent node—not merely its removal. Upward change to a parental hub node from its children must include the complication of its priors.

Eigenvector centrality and its variants (Katz centrality and PageRank, for example) appear to be the primary candidates for measures of ‘most influential node’ in networks far simpler than those envisaged here, despite a number of critical studies (Borgatti and Everett, 2006; Newman, 2018; Da Silva et al., 2012; Chen et al., 2013; Dablander and Hinne, 2019). But at their best these too are most appropriate for undirected networks modeling something like infection dynamics or directed networks modeling something like internet searches. They count transitive numbers of contacts, but don’t include the complexities of prior credences, conditional probabilities, and the likelihood ratios of evidence that are an inherent part of change within epistemic networks.

We can analyze some of the complexities involved, working towards general rules of thumb for epistemic impact, by considering three simple networks. Figure 4, considered above, has the structure of a binary tree. We represent its structure alone in the second image of Figure 7. The binary tree can be thought of as a structure intermediate between two others: a purely linear network, shown in the first image, and hub or ‘star’ shown in the third image in Figure. 7

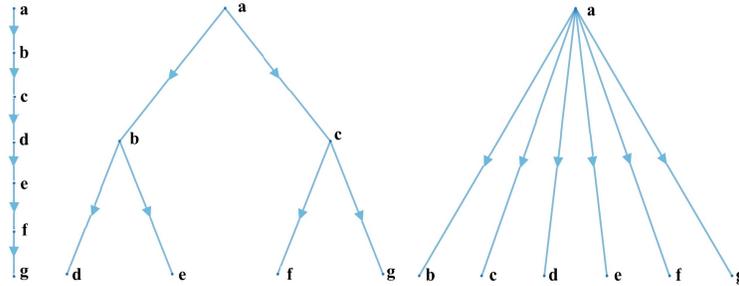


Figure 7: Three basic networks: linear, binary tree, and hub or star

The binary tree may represent a theoretical structure in which a fundamental theory *a* supports two derivative theories *b* and *c* independently, each of which supports two further hypotheses. The purely linear hypothesis exaggerates the vertical descent of such a structure—a unary ‘tree’—in which a fundamental theory supports a derivative theory at *b*, supporting a hypothesis at *c* which in turn supports a hypothesis at *d*. One can think of the hub or star figure as exaggerating the horizontal spread within a theory. Here, a single fundamental theory supports directly subsidiary hypotheses at *b*, *c*, and so forth. These, then, are differences in over-all network structure. Within those structures our Kullback-Leibler and Brier divergence graphs show the network impact at different nodes of evidence across the range of likelihood ratios. We can further consider variations in prior credences and conditional probabilities. We will emphasize Brier divergence simply because of its conceptual simplicity and greater familiarity in the philosophical literature.

For purposes of the rough rules of thumb we aim for here, we can think of a credence as high if it is greater than .5 (corresponding to a greater probability that it is true than that it is false), and as low if it is less than 0.5. In line with the exponents on our x-axes, we can think of evidence with a likelihood ratio greater than 1 as positive, with evidence with a likelihood ratio less than 1 as negative. Negative evidence in the case of a hypothesis with low credence, like positive evidence in the case of a hypothesis with high credence, we will term ‘credence-reinforcing’ or simply ‘reinforcing’ regarding the current credence. The opposite we will term ‘credence-counter’ or simply ‘counter evidence.’⁹ Finally, we can

9. We’ve resisted a temptation to use the term ‘confirmatory’ for negative evidence in the case of low credences, and positive evidence in the case of high credences, in order to avoid confusion with standard terminology in which ‘*e* confirms *h*’ is used to indicate that *e* raises the probability of *h* regardless of *h*’s initial credence. We are obliged to an anonymous referee for pointing out the potential confusion.

think of a pair of conditional probabilities $p(y | x)$ and $p(y | \sim x)$ as positive if $p(y | x) > 0.5$ and $p(y | \sim x) < 0.5$, as negative if $p(y | x) < 0.5$ and $p(y | \sim x) > 0.5$. Here we consider only symmetrical conditional probabilities, such that $p(y | x)$ and $p(y | \sim x)$ sum to 1, and use the same conditional probabilities on all links. These are particularly severe limitations in this first exploration of such a model; we start with them here merely for the sake of simplicity.

The simplest case of our three examples, though illustrative of some general principles that hold throughout, is the hub or star of figure 7. Figures 8a and 8b show the case in which conditional probabilities are positive throughout (uniformly [0.3, 0.7]) and our root node is given a high value of 0.6 on the left and a low 0.4 on the right, with derivative high credences of 0.54 for all other nodes on the left and a low 0.46 on the right.

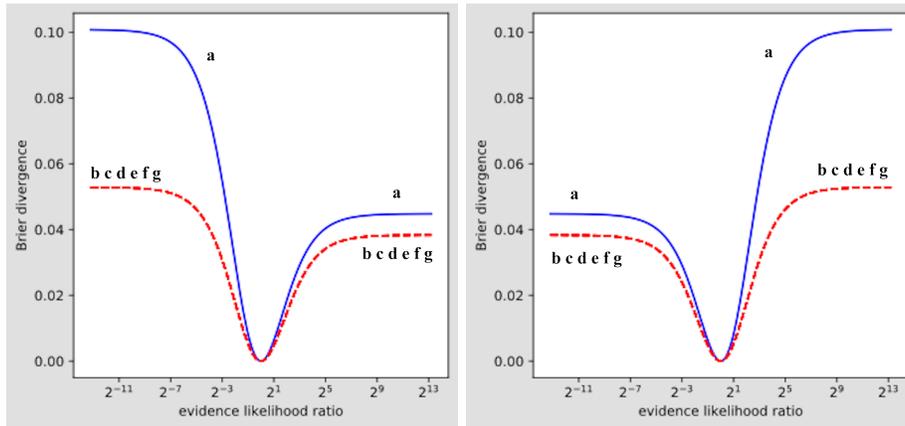


Figure 8: Brier divergence at different likelihood ratios for introduced evidence in the star network, with positive conditional probabilities ([0.3, 0.7]) and priors of 0.6 and 0.54 at root and leaf nodes (8a) and 0.4 and 0.46 at root and leaf nodes (8b)

The clear left-right reversal of the two graphs follows a simple pattern that holds throughout the examples we give not only because the priors of 0.6 and 0.4 that we use as examples are symmetrical around .5 but because of the assumption noted that our conditional probabilities for $p(y | x)$ and $p(y | \sim x)$ sum to 1. Both of these are assumptions are made simply for the sake of simplicity in exposition. Varying either of these modeling assumptions breaks the symmetry shown.

The fact that values are higher on the left in 8a and on the right in 8b follows simple Bayesian principles. We offer figure 9 as a reminder of the differential impact of evidence e gauged in terms of likelihood ratio contingent on the prior for a hypothesis h.

A high prior given evidence with what we've termed a positive likelihood

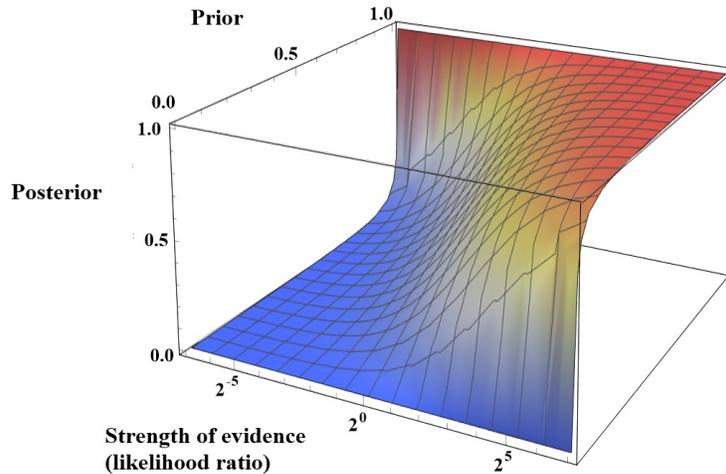


Figure 9: A reminder of the immediate Bayesian impact of evidence at a node, in terms of likelihood ratio, depending on its prior

ratio (> 1) leaves us with a high posterior credence, but our posterior plummets sharply as the likelihood ratio declines and turns negative. A low prior given evidence with a negative likelihood ratio (< 1) leaves us with a low posterior, but that posterior climbs sharply as the likelihood ratio climbs and turns positive. With the conditional probabilities in place in this case that effect carries through our example network as a whole. In 8a our credences for all nodes are high (> 0.5), with the result that positive evidence on the right side of the graph produces a lesser impact than does negative evidence on the left. In 8b all credences are symmetrically low (< 0.5), explaining the reverse pattern.

In this case illustrated in Figure 8, regardless of whether evidence is positive or negative, it is the root node that is the most influential. With this network, these conditional probabilities and these priors, it is uniformly credence change at the root that has the greatest impact. That changes, however, if our conditional probabilities are changed from positive to negative in our rough sense, from $[0.3, 0.7]$ to $[0.7, 0.3]$ —from probabilities $p(y | x)$ of 0.7 and $p(y | \sim x)$ of 0.3 for a node y and its parent x to probabilities $p(y | x)$ of 0.3 and $p(y | \sim x)$ of 0.7. With a prior of 0.6 at node a priors for our leaf nodes are 0.46 in this case, giving us the Brier divergence graphs of Figure 10.

The reversal in our two graphs is the same as before. The fact that our root node has a higher impact than leaf nodes in the case of negative evidence in 10a, but lower in the case of positive evidence, is explained by the disparity in our priors. The prior for our root node is high; that for our leaf nodes is low. In this

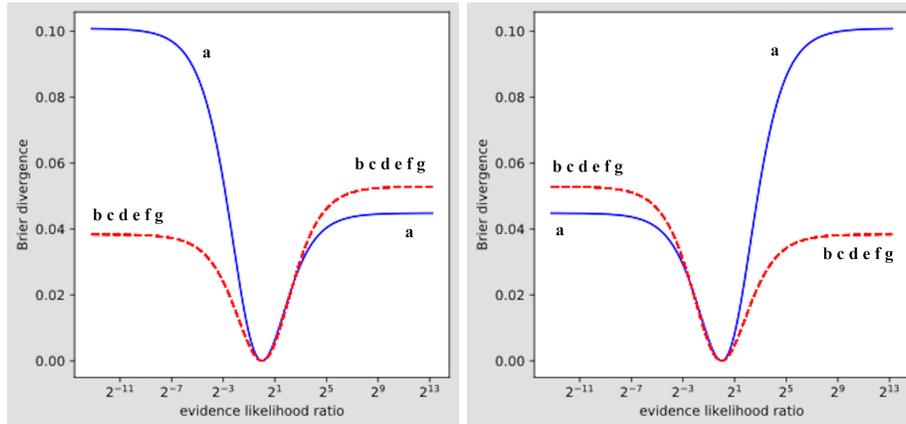


Figure 10: Brier divergence at different likelihood ratios for introduced evidence in the star network, with positive conditional probabilities ([0.7, 0.3]) and priors of 0.6 and 0.46 at root and leaf nodes (10a) and 0.4 and 0.54 at root and leaf nodes (10b)

case, negative evidence constitutes counter-evidence for node a while reinforcing credences in the leaf nodes, resulting in a higher credence change for the former than the latter, here reflected in changes in the network as a whole. Positive evidence, on the other hand, will be counter-evidence for our leaf nodes while reinforcing for our root node, reflected in the reversal of effects from the left to the right side of 10a. A mirror-image explanation applies for the mirror image graph of 10b.

The interactive factors of priors, conditional probabilities, and evidence likelihood play out along similar lines but with different effects in the different structures of our other sample networks. Figures 11a and 11b show Brier divergence in a binary tree structure with priors set at 0.6 (11a) and 0.4 (11b) and with negative conditional probabilities of [0.7, 0.3].

Here, the pattern is similar to that analyzed in the preceding section. For negative evidence in the case of 11a, it is the root node that has the greatest influence, followed by the central nodes and finally the leaf nodes. This is what one might expect from the same evidence impact at nodes with descending high priors. For positive evidence, in contrast, it is the central nodes b and c that have the most influence, followed by the leaf nodes and finally by the root. Here secondary effects of updating through the network dominate the simple effect of evidence at impact nodes.

Figures 12a and 12b show Brier divergence with the same binary tree structure and priors but in which conditional probabilities are set at [0.7, 0.3] throughout

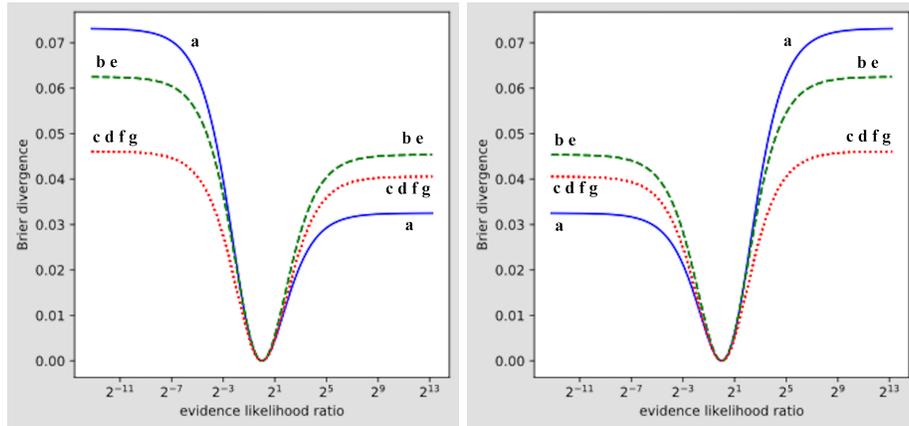


Figure 11: Brier divergence at different likelihood ratios for introduced evidence in the binary tree network, with negative conditional probabilities [0.7, 0.3]. In (a), a root node prior of 0.6 at node a entails credences of 0.54 at central branch nodes and 0.516 at leaves. In (b), a root node of 0.4 entails credences of 0.46 at central branch nodes and 0.484 at leaves.

instead of [0.3, 0.7].

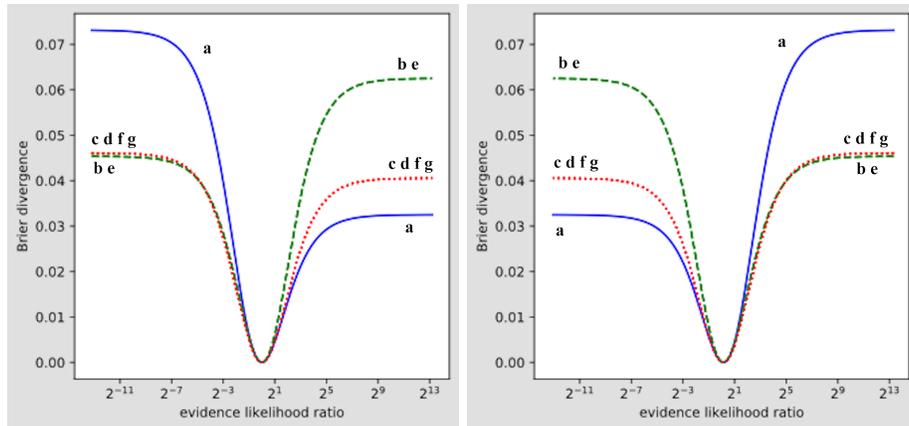


Figure 12: Brier divergence at different likelihood ratios for introduced evidence in the binary tree network, with positive conditional probabilities [0.3, 0.7]. In (a), a root node prior of 0.6 at node a entails credences of 0.46 at central branch nodes and 0.516 at leaves. In (b), a root node of 0.4 entails credences of 0.54 at central branch nodes and 0.484 at leaves.

The role of positive evidence in the case of 12a and of negative evidence in the case of 12b follows the same pattern as in the case of reversed conditional

probabilities in 11a and b. The difference in conditional probabilities makes a clear difference on the other sides of the graphs. In the case of negative evidence in 12a and positive evidence in 12b, the change in conditional probabilities collapses the impact of central and leaf nodes. Here evidence at the root node has a high impact, with evidence at all other nodes effectively the same, roughly as in the case of the star.

Of our three basic network structures, it is perhaps surprising that a simple linear network proves to be the most complex. Figures 13a and 13b show graphs of Brier divergence for the purely linear network of Figure 7 where conditional probabilities are again set uniformly at a positive [0.3, 0.7]. In 13a a root node initiated at 0.6 gives us descending credences 0.54, 0.516, 0.5064, 0.50256, 0.501024, and 0.50041 at nodes *b* through *g*. In 13b a root node of 0.4 gives us descending credences of 0.4, 0.46, 0.484, 0.493, 0.497, 0.498, and 0.49959.

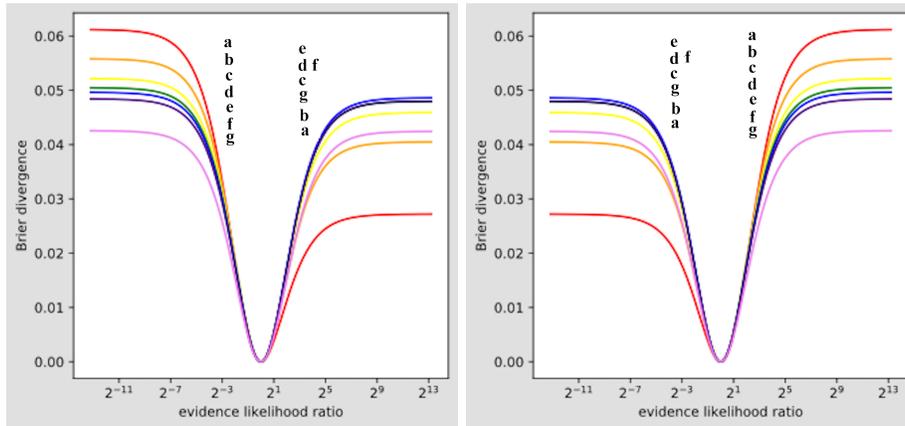


Figure 13: Brier divergence at different likelihood ratios for introduced evidence in the binary tree network, with positive conditional probabilities [0.3, 0.7]. In (a), a root node prior of 0.6 at node *a* entails credences of 0.46 at central branch nodes and 0.516 at leaves. In (b), a root node of 0.4 entails credences of 0.54 at central branch nodes and 0.484 at leaves.

Basic principles noted above still apply: our exploration of symmetrical priors with an assumption of conditional probabilities that add to 1 still dictate a mirror symmetry between the two cases, and counter-evidence in our sense will have a greater impact on credence change in a node than will reinforcing evidence. But the pattern does not show a simple descent in over-all node influence in terms of either descending network order or diminishing prior credence (in the case of 13a) or increasing prior credence (in the case of 13b). The fact that evidence at node *a* has the highest impact in counter-evidence cases and the lowest in credence-

reinforcing cases may be credited to its relatively high prior credence. But the fact that evidence at node b, with the second highest prior, shows the lowest impact of all nodes in counter-evidence cases and the highest in credence-reinforcing cases calls for an explanation that includes node position in the network and the asymmetrical strength of upward influence through conditioning, which is sensitive to the priors of upwardly influenced nodes, and downward influence through our conditional probabilities, which is not sensitive to downstream priors. In each case those asymmetrical effects are magnified by the transitive effects through successive nodes.

The situation becomes still more complex when our conditional probabilities are set at $[0.7, 0.3]$ instead of $[0.3, 0.7]$, shown in Figure 14 for root node priors of 0.6 and 0.4.

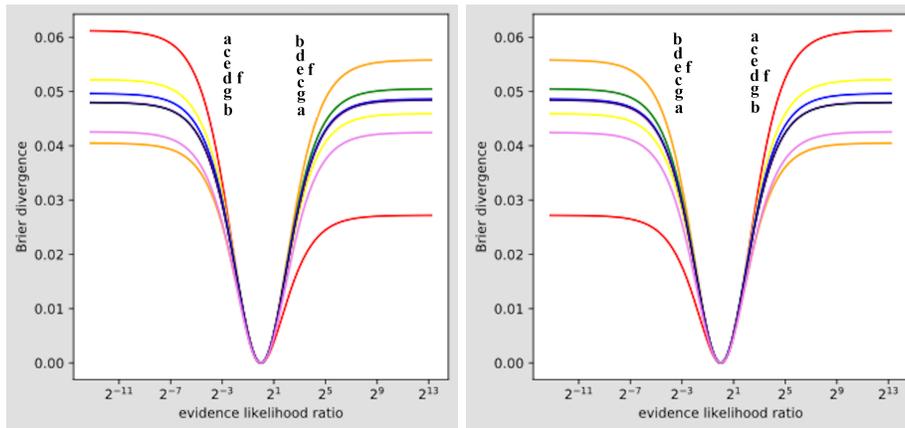


Figure 14: Brier divergence at different likelihood ratios for introduced evidence in the linear network, with negative conditional probabilities $[0.7, 0.3]$. In (a), a root node prior of 0.6 at node a entails descending credences of 0.46, 0.516, 0.4936, 0.50256, 0.498976, and 0.50041 for nodes b through g. In (b), a root node of 0.4 entails descending credences of 0.54, 0.484, 0.5064, 0.49744, 0.501024 and 0.49959.

Here, unlike in the preceding cases, our priors are neither all high (>0.5) nor all low (<0.5). Evidence with the same likelihood ratio will therefore affect different nodes differently. Upward transfer by conditioning, moreover, will be differently affected by the differently high or low credences of nodes through which it proceeds.

The variety of effects in even these simple structures, with differently positive or negative evidence, conditional probabilities, and high or low initial priors, offer some hints toward the form that a network metric for evidence influence in theo-

retical networks would have to take. It is clear that influence metrics appropriate to the simple contact networks that have been studied, directed or undirected—without node credences, without differential evidence characteristics, and without link complexities of conditional probabilities and the complexities of Bayesian updating and Jeffrey conditionalization—will be inadequate here. Development of a Bayesian net analogue to eigenvector centrality that includes these factors seems a promising if forbidding prospect, but is beyond the scope of our work here.

5 Theoretical Structures Beyond Polytrees

In work to this point, for the sake of simplicity, we have concentrated on polytrees: structures in which each node has a single parent. Things become more complex when we expand consideration—as we should—to the wider range of directed acyclic graphs in general. Both of our initial examples—the causal structure of the failure of the 17th street levee in New Orleans and the grounding structure of theory regarding the COVID-19 pandemic—include not merely branches but downward junctures or ‘colliders’ from multiple nodes as well as branches, well beyond simple polytrees.

Figure 15 shows the simplest downward juncture, in which credence for c is determined by values for both a and b . The simple conditional credences used above must be replaced with a conditional probability matrix, reflecting values for c contingent on combinatory values for a and b . Figure 16 shows Brier divergence for a case in which we start with a credence of .6 at both a and b and in which our matrix is ‘or-like’, specifying a value for c of 0.9 if either a or b is true, a value of 0.1 otherwise. Figure 17 shows Brier divergence for the same values at a and b but a conditional matrix that is ‘and-like’, specifying a value of 0.9 for c only if both a and b are true, a value of 0.1 otherwise.

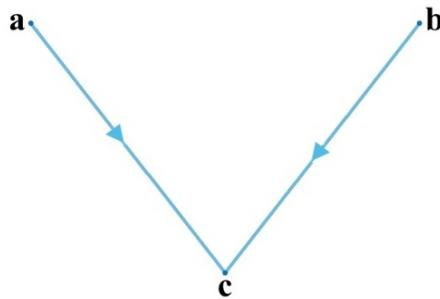


Figure 15: A simple downward juncture

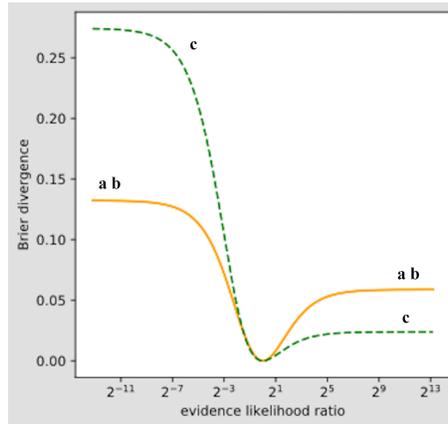


Figure 16: Brier divergence at different likelihood ratios for introduced evidence in a simple ‘or-like’ juncture, with conditional probabilities [0.1, 0.9] for all cases except $a = 0$ and $b = 0$, where conditional probabilities are [0.9, 0.1]. Root node priors of 0.6 at both a and b entail a prior credence of 0.772 for c .

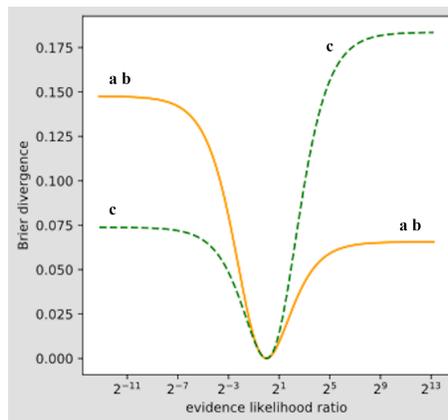


Figure 17: Brier divergence at different likelihood ratios for introduced evidence in a simple ‘and-like’ juncture, with conditional probabilities [0.9, 0.1] for all cases except $a = 1$ and $b = 1$, where conditional probabilities are [.1, .9]. Root node priors of 0.6 at both a and b entail a prior credence of 0.338 for c .

With the same priors at root nodes and conditional probability ratios of 0.1 and 0.9, the difference between ‘and-like’ and ‘or-like’ junctures is clear even in the priors entailed for the juncture node c : a positive 0.772 in the case of an ‘or-like’ juncture, 0.338 in the case of ‘and’. The result is a reversal of evidence importance at nodes: the sensitivity of c dominates that of a and b with negative evidence in

the case of an ‘or’ juncture, while both a and b dominate c with negative evidence in the case of ‘and.’ Appropriately to their status as duals, the cases are reversed for positive evidence.

We have already noted the quick complexity of evidence impact contingent on theoretical structure, priors, and conditional probabilities in the case of polytrees. Expansion to full directed acyclic graphs increases that complexity significantly even with the simple ‘and-like’ and ‘or-like’ junctures illustrated, let alone for the far richer palette of possible probability distributions.

The tools outlined do offer us greater prospects for understanding the parameters of evidence impact in different theoretical structures, however. As a final illustration, while fully recognizing its limitations, we offer a Brier divergence graph for the COVID-19 graph with which we began (Figure 18). We use the formal treatment outlined in previous section, which builds in independence assumptions that are undoubtedly unrealistic. Full specifications for our largely ad hoc assignments of priors and conditional probabilities in this case are documented in the Appendix.

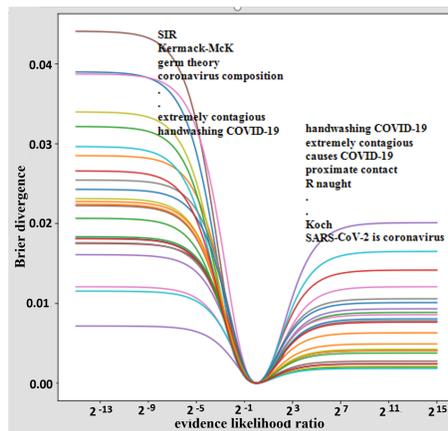


Figure 18: Evidence sensitivity in the theoretical structure for COVID-19 shown in Figure 2. Full specifications for conditional probabilities and priors, as well as Brier scores for all nodes, appear in the Appendix.

It is perhaps not surprising that the theoretical structure exhibited for the coronavirus is most sensitive to negative or counter-evidence, which appears on the left side of the graph, at or near the root nodes. Were we to find out that our basic assumptions regarding germ theory, regarding the susceptible-infected-recovery SIR model, or composition of viruses were incorrect, our credence in the other propositions in the theory would be importantly impacted. On the right side

of the graph, it is perhaps more surprising that positive or credence-reinforcing evidence most affects the graph with confirmation that handwashing stops the spread of COVID-19. It is perhaps less surprising that confirmation that SARS-CoV-2 causes COVID-19 strongly strengthens credence in other elements of the theoretical structure. As detailed in the Appendix, however, there is by no means a simple inversion in node order with regard to evidence impact from the left to the right side of the graph. Evidence regarding an R_0 for the Coronavirus, for example, appears relatively high on both the left and the right side of the graph, in 8th and 5th places respectively.

6 Conclusion

The structure of scientific theories can be modeled as Bayesian nets, with values at nodes modeling credence in different propositions in the structure and conditional probabilities modeling the inferential and evidential connections between elements of the theory. Precisely because of that structure, a scientific theory will be more vulnerable to the impact of evidence of the same strength at different nodes, modeled in terms of the Bayesian impact of evidence of a particular likelihood ratio at that node.

This conclusion has two clear epistemic implications, both of which are probably intuitive but both of which are captured here in a quantitative model.

The first implication is that relative position of different nodes within a single theoretical structure can make a major epistemic difference. Given a single theoretical structure, evidence impact at one point in a structure will have a different effect on credences, percolated through the network, than will evidence impact at another point. Calculating the network factors that lead to greater influence, however, is far from simple, depending not merely on the skeletal structure of a network but on the specific conditional probabilities on inferential and evidential links, the specific prior credences within the network, and the likelihood-ratio character of the evidence itself.

The second and related implication is that the differences between different theoretical structures can make a major epistemic difference. The same piece of evidence at the leaf node of a star-like network will have an importantly different effect than at the leaf node of a linear network, or a binary tree. Here again, however, calculating the aspects of theoretical structure that make the difference in evidence impact in different cases proves to be far from simple.

These are first steps; we have emphasized the preliminary and suggestive character of our work throughout. The promise of an approach to scientific theories

as Bayesian nets is the promise of being better able to understand the theoretical-structure-relative dynamics of scientific evidence and scientific change. Fulfilling that promise will demand development of a more complete understanding of node-importance centrality in Bayesian nets, further modeling of dynamic change within Bayesian nets with iterated credence changes, and modeling of structural changes in the theoretical structures themselves. That is where some of the questions of greatest interest lie, but all of that we leave to further work.

Appendix

Specifications of root node priors and conditional probabilities for sensitivity calculations in the reconstruction of the covid-19 theoretical structure shown in section 7.

Nodes:

a = germ theory
b = transmission
c = Koch
d = causes COVID-19
e = respiratory infection
f = airborne and contact
g = proximate human contact
h = social distancing
i = coronavirus composition
j = SARS-CoV-2 is coronavirus
k = breaking lipid
l = soap and water
m = handwashing coronaviruses
n = handwashing SARS-CoV-2
o = handwashing COVID-19
p = SIR
q = Kermack-McKendrick
r = R naught
s = COVID-19 R_0 1.4 – 2.0
t = extremely contagious
u = sigmoid curve
v = levelling off

w = acquired immunity
 x = COVID-19 extinction
 Root nodes a, c, i, j , and p : prior credences 0.8.
 Nodes l, f, s, w : prior credences 0.7

Conditional probabilities:

$p(b | a) = 0.8; p(b | \sim a) = 0.2$
 $p(d | b, c) = 0.8; 0.2$ otherwise
 $p(e | d, f) = 0.8; p(e | d, \sim f) = 0.5; p(e | \sim d, f) = 0.5; p(e | \sim d, \sim f) = 0.2$
 $p(g | f) = 0.8; p(g | \sim f) = 0.2$
 $p(h | g) = 0.8; p(h | \sim g) = 0.4$
 $p(k | i) = 0.7; p(k | \sim i) = 0.2$
 $p(m | l, k) = 0.8; p(m | \sim k, l) = 0.5; p(m | \sim l, k) = 0.5; p(m | \sim l, \sim k) = 0.2$
 $p(n | m, j) = 0.8; p(n | m, \sim j) = 0.5; p(n | \sim m, j) = 0.5; p(n | \sim m, \sim j) = 0.2$
 $p(o | d, n) = 0.8; 0.2$ otherwise
 $p(q | p) = 0.8; p(q | \sim p) = 0.2$
 $p(r | q) = 0.8; p(r | \sim q) = 0.2$
 $p(u | q) = 0.8; p(u | \sim q) = 0.2$
 $p(t | r, s) = 0.8; 0.2$ otherwise
 $p(v | u) = 0.8; p(v | \sim u) = 0.5$
 $p(x | j, d, w, v) = 0.9; p(x | \sim j, d, w, v) = 0.8; p(x | j, \sim d, w, v) = 0.8; p(x | j, \sim d, w, v) = 0.7; p(x | j, d, \sim w, v) = 0.8; p(x | \sim j, d, \sim w, v) = 0.7; p(x | j, \sim d, \sim w, v) = 0.6; p(x | \sim j, d, w, \sim v) = 0.8; p(x | \sim j, d, w, \sim v) = 0.7; p(x | j, \sim d, w, \sim v) = 0.6; p(x | j, d, \sim w, \sim v) = 0.7; p(x | \sim j, d, \sim w, \sim v) = 0.6; p(x | j, \sim d, \sim w, \sim v) = 0.6; p(x | \sim j, \sim d, \sim w, \sim v) = .02$

Brier divergence scores by node for evidence likelihood 2^{-5} , left side of figure 21:

0.032702925196386255 p SIR
 0.03197274586977256 q Kermack-McKendrick
 0.02892300064134992 a germ theory
 0.025206021616052884 i coronavirus composition
 0.023849404607912673 c Koch
 0.023530107955275256 b transmission
 0.021985716054253094 j SARS-CoV-2 is coronavirus
 0.021732855368419703 r R naught

0.020731908792273795 *u* sigmoid curve
 0.02050306150659843 *x* COVID-19 extinction
 0.018859774335903478 *s* COVID-19 R_0 1.4
 0.01876352205497495 *v* levelling off
 0.018291007460366212 *l* soap and water
 0.018142736120275237 *f* airborne and contact
 0.01684255830020106 *w* acquired immunity
 0.01577183852380332 *m* handwashing coronaviruses
 0.015542370861284812 *k* breaking lipid
 0.015403295972751313 *n* handwashing SARS-CoV-2
 0.0153470017407081 *d* causes COVID-19
 0.015074719685156309 *h* social distancing
 0.013945521086314742 *e* respiratory infection
 0.01066307009488828 *g* proximate contact
 0.01028733771960207 *t* extremely contagious
 0.006463834532320453 *o* handwashing COVID-19

Brier divergence scores by node for evidence likelihood 2^5 , right side of figure 21:

0.01664743806591248 *o* handwashing COVID-19
 0.01439930691393785 *t* extremely contagious
 0.012583132032143445 *d* causes COVID-19
 0.010663070094888275 *g* proximate contact
 0.009542603689210295 *r* R naught
 0.009103101547024603 *u* sigmoid curve
 0.008347243506501229 *e* respiratory infection
 0.007966962560939319 *m* handwashing coronaviruses
 0.007820365918820486 *q* Kermack-McKendrick
 0.0072646504369333125 *k* breaking lipid
 0.007046065875329723 *h* social distancing
 0.006301081627599784 *n* handwashing SARS-CoV-2
 0.005755340972874356 *b* transmission
 0.004496778986033397 *v* levelling off
 0.0038829270106076202 *s* COVID-19 R_0 1.4
 0.003765826974073341 *l* soap and water
 0.0037353002678103607 *f* airborne and contact
 0.003467614372621885 *w* acquired immunity
 0.002546901691876484 *p* SIR

0.0022525214128471527 *a* germ theory
0.002231827418619689 *x* COVID-19 extinction
0.001963043327588758 *i* coronavirus composition
0.0018573900830391689 *c* Koch
0.0017122461394334491 *j* SARS-CoV-2 is coronavirus

References

- Aguilar, J. (2013). Different dynamic causal relationship approaches for cognitive maps. *Applied Soft Computing*, 13(1):271–282.
- Albert, R., Jeong, H., and Barabási, A.-L. (2000). Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382.
- Alexander, J. M., Himmelreich, J., and Thompson, C. (2015). Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science*, 82(3):424–453.
- Andersen, C., Battjes, J., Daniel, D., Edge, B., and Espey, W. (2007). American society of civil engineers. hurricane katrina external review panel. *The New Orleans hurricane protection system: what went wrong and why: a report*.
- Anderson, R. M. and May, R. M. (1979). Population biology of infectious diseases: Part i. *Nature*, 280(5721):361–367.
- Axelrod, R. (1976). The cognitive mapping approach to decision making. *Structure of decision*, pages 221–250.
- Bandyopadhyay, P. S., Brittan, G. G., and Taper, M. L. (2016). *Belief, evidence, and uncertainty: problems of epistemic inference*. Springer.
- Bao, Z.-K., Ma, C., Xiang, B.-B., and Zhang, H.-F. (2017). Identification of influential nodes in complex networks: Method from spreading probability viewpoint. *Physica A: Statistical Mechanics and its Applications*, 468:391–397.
- Bea, R. G. (2008). Failure of the new orleans 17th street canal levee and floodwall during hurricane katrina. In *From Research to Practice in Geotechnical Engineering*, pages 574–593.
- Borgatti, S. P. and Everett, M. G. (2006). A graph-theoretic perspective on centrality. *Social networks*, 28(4):466–484.
- Bovens, L. and Hartmann, S. (2002). Bayesian networks and the problem of unreliable instruments. *Philosophy of Science*, 69(1):29–72.
- Bovens, L., Hartmann, S., et al. (2003). *Bayesian epistemology*. Oxford University Press on Demand.
- Boyd, E. C. (2010). Estimating and mapping the direct flood fatality rate for flooding in greater new orleans due to hurricane katrina. *Risk, Hazards & Crisis in Public Policy*, 1(3):91–114.

- Cakmak, E. H., Dudu, H., Eruygur, O., Ger, M., Onurlu, S., and Tonguç, Ö. (2013). Participatory fuzzy cognitive mapping analysis to evaluate the future of water in the seyhan basin. *Journal of water and climate change*, 4(2):131–145.
- Champion, C. and Elkan, C. (MS). Visualizing the consequences of evidence in bayesian networks. *arXiv preprint arXiv:1707.00791*.
- Chen, D.-B., Gao, H., Lü, L., and Zhou, T. (2013). Identifying influential nodes in large-scale directed networks: the role of clustering. *PloS one*, 8(10):e77455.
- Climenhaga, N. (2019). The structure of epistemic probabilities. *Philosophical Studies*, pages 1–30.
- Climenhaga, N. (2020). Evidence and inductive inference. In Littlejohn, C. and Lasonen-Aarnio, M., editors, *The Routledge Handbook of the Philosophy of Evidence*, ed. Maria LasonenAarnio and Clayton Littlejohn. Routledge.
- Crupi, V. (2020). Confirmation. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2020 edition.
- Crupi, V., Tentori, K., and Gonzalez, M. (2007). On bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science*, 74(2):229–252.
- Da Silva, R. A. P., Viana, M. P., and da Fontoura Costa, L. (2012). Predicting epidemic outbreak from individual features of the spreaders. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(07):P07005.
- Dablander, F. and Hinne, M. (2019). Node centrality measures are a poor substitute for causal inference. *Scientific Reports (Nature Publisher Group)*, 9(1).
- Easwaran, K. (2011). Bayesianism i: Introduction and arguments in favor. *Philosophy Compass*, 6(5):312–320.
- Findlay, S. D. and Thagard, P. (2014). Emotional change in international negotiation: Analyzing the camp david accords using cognitive-affective maps. *Group Decision and Negotiation*, 23(6):1281–1300.
- Goodman, N. (1955). *Fact, Fiction, and Forecast*. Harvard University Press.
- Hafiene, N., Karoui, W., and Romdhane, L. B. (2019). Influential nodes detection in dynamic social networks. In *International Conference on Business Information Systems*, pages 62–73. Springer.
- Hartmann, S. (2008). Modeling in philosophy of science. *Representation, evidence, and justification: Themes from Suppes*, 1:95–121.

- Hassan, S. A., Sheikh, F. N., Jamal, S., Ezeh, J. K., and Akhtar, A. (2020). Coronavirus (covid-19): a review of clinical features, diagnosis, and treatment. *Cureus*, 12(3).
- Henderson, L., Goodman, N. D., Tenenbaum, J. B., and Woodward, J. F. (2010). The structure and dynamics of scientific theories: A hierarchical bayesian perspective. *Philosophy of Science*, 77(2):172–200.
- Hobbs, B. F., Ludsin, S. A., Knight, R. L., Ryan, P. A., Biberhofer, J., and Ciborowski, J. J. (2002). Fuzzy cognitive mapping as a tool to define management objectives for complex ecosystems. *Ecological Applications*, 12(5):1548–1565.
- Homer-Dixon, T., Milkoreit, M., Mock, S. J., Schröder, T., and Thagard, P. (2014). The conceptual structure of social disputes: Cognitive-affective maps as a tool for conflict analysis and resolution. *SAGE Open*, 4(1):2158244014526210.
- Jeffrey, R. C. (1965). *The Logic of Decision*. 1st Edition, University of Chicago Press.
- Jetter, A. J. and Sperry, R. C. (2013). Fuzzy cognitive maps for product planning: using stakeholder knowledge to achieve corporate responsibility. In *2013 46th Hawaii International Conference on System Sciences*, pages 925–934. IEEE.
- Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 263–297. Synthese.
- Kempe, D., Kleinberg, J., and Tardos, É. (2005). Influential nodes in a diffusion model for social networks. In *International Colloquium on Automata, Languages, and Programming*, pages 1127–1138. Springer.
- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721.
- Kosko, B. (1986). Fuzzy cognitive maps. *International journal of man-machine studies*, 24(1):65–75.
- Kuhn, T. S. (2012 [1969]). *The structure of scientific revolutions*. University of Chicago press.
- Lakatos, I. (1969). II—criticism and the methodology of scientific research programmes. *Proceedings of the Aristotelian Society*, 69(1):149–186.

- Laudan, L. (1978). *Progress and its problems: Towards a theory of scientific growth*, volume 282. Univ of California Press.
- Lawyer, G. (2015). Understanding the influence of all nodes in a network. *Scientific reports*, 5(1):1–9.
- Li, M., Zhang, Q., and Deng, Y. (2018). Evidential identification of influential nodes in network of networks. *Chaos, Solitons & Fractals*, 117:283–296.
- Newman, M. (2018). *Networks*. Oxford university press.
- Olsson, E. J. (2013). A bayesian simulation model of group deliberation and polarization. In *Bayesian argumentation*, pages 113–133. Springer.
- Olsson, E. J. and Vallinder, A. (2013). Norms of assertion and communication in social networks. *Synthese*, 190(13):2557–2571.
- Papageorgiou, E. and Stylios, C. (2008). Fuzzy cognitive maps. *Handbook of granular computing*, 123:755–775.
- Pearl, J. (1988). *Probabilistic Inference in Intelligent Systems*. San Matteo. CA: Morgan Kaufmann.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press UK.
- Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review*, 60(1):20–43.
- Quine, W. V. O. and Ullian, J. S. (1970). *The Web of Belief*. New York: Random House.
- Rogers, J. D., Boutwell, G. P., Schmitz, D. W., Karadeniz, D., Watkins, C. M., Athanasopoulos-Zekkos, A., and Cobos-Roa, D. (2008). Geologic conditions underlying the 2005 17th street canal levee failure in new orleans. *Journal of geotechnical and geoenvironmental engineering*, 134(5):583–601.
- Schaffer, J. (2016). Grounding in the image of causation. *Philosophical studies*, 173(1):49–100.
- Shafer, G. (1981). Jeffrey’s rule of conditioning. *Philosophy of Science*, 48(3):337–362.
- Skyrms, B. (1984). *Pragmatics and Empiricism*. Yale University Press, New Haven.

- Skyrms, B. (1987). Coherence. *Scientific inquiry in philosophical perspective*, pages 225–242.
- Skyrms, B. and Lambert, K. (1995). The middle ground: Resiliency and laws in the web of belief. *Laws of Nature: Essays on the Philosophical, Scientific and Historical Dimensions*. de Gruyter.
- Soler, L. S., Kok, K., Camara, G., and Veldkamp, A. (2012). Using fuzzy cognitive maps to describe current system dynamics and develop land cover scenarios: a case study in the brazilian amazon. *Journal of Land Use Science*, 7(2):149–175.
- Spirtes, P. (2010). Introduction to causal inference. *Journal of Machine Learning Research*, 11(5).
- Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. (2000). *Causation, prediction, and search*. MIT press.
- Spohn, W. (2012). *The laws of belief: Ranking theory and its philosophical applications*. Oxford University Press.
- Thagard, P. (2015). The cognitive–affective structure of political ideologies. In *Emotion in group decision and negotiation*, pages 51–71. Springer.
- van Vliet, M., Kok, K., and Veldkamp, T. (2010). Linking stakeholders and modellers in scenario studies: The use of fuzzy cognitive maps as a communication and learning tool. *Futures*, 42(1):1–14.
- Wei, H., Pan, Z., Hu, G., Zhang, L., Yang, H., Li, X., and Zhou, X. (2018). Identifying influential nodes based on network representation learning in complex networks. *PloS one*, 13(7):e0200091.