Synthese (2022) 200:297
https://doi.org/10.1007/s11229-022-03720-z

ORIGINAL RESEARCH

The punctuated equilibrium of scientific change: a Bayesian network model

Patrick Grim[1,2] · Frank Seidl[3] · Calum McNamara[4] · Isabell N. Astor[5] · Caroline Diaso[3]

Received: 16 April 2021 / Accepted: 26 April 2022 / Published online: 12 July 2022

Abstract
Our scientific theories, like our cognitive structures in general, consist of propositions linked by evidential, explanatory, probabilistic, and logical connections. Those theoretical webs 'impinge on the world at their edges,' subject to a continuing barrage of incoming evidence (Quine 1951, 1953). Our credences in the various elements of those structures change in response to that continuing barrage of evidence, as do the perceived connections between them. Here we model scientific theories as Bayesian nets, with credences at nodes and conditional links between them modelled as conditional probabilities. We update those networks, in terms of both credences at nodes and conditional probabilities at links, through a temporal barrage of random incoming evidence. Robust patterns of punctuated equilibrium, suggestive of 'normal science' alternating with 'paradigm shifts,' emerge prominently in that change dynamics. The

✉ Patrick Grim
patrick.grim@stonybrook.edu

Frank Seidl
fcseidl@umich.edu

Calum McNamara
camcnam@umich.edu

Isabell N. Astor
astori@umich.edu

Caroline Diaso
cdiaso@umich.edu

1 Philosophy, Stony Brook University, Stony Brook, NY 11794, USA

2 Center for Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109, USA

3 Mathematics, University of Michigan, Ann Arbor, MI 48109, USA

4 Philosophy, University of Michigan, Ann Arbor, MI 48109, USA

5 Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

suggestion is that at least some of the phenomena at the core of the Kuhnian tradition are predictable in the typical dynamics of scientific theory change captured as Bayesian nets under even a random evidence barrage.

Keywords Bayesian networks · Scientific change · Paradigm shifts · Punctuated equilibrium

1 Webs of belief and the impact of evidence

The model we explore here grows from suggestive remarks in Quine (1951, 1953), further refined in Skyrms and Lambert (1995):

> 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….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… (Quine 1953, p. 42)

> But however attractive this picture may be, Quine does not offer any methodology for modeling and mapping the networks of belief….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. (Skyrms & Lambert, 1995, pp. 139–140).

We model Quine's webs of belief as Bayesian networks, very much along the lines that Skyrms and Lambert suggest. Nodes represent propositions, carrying specific credences or degrees of belief. Connections of implication and support within a theoretical structure are modeled by links that carry conditional probabilities. Within such a network, evidence at one node results in a cascade of credence changes, large or small, elsewhere in the network.[1] We tend to think of theoretical networks as abstract objects composed of linked propositions, avoiding controversies pitting credences against 'full' beliefs, but the basic structural picture is the same as that invoked by Quine, Skyrms and Lambert.

Both Quine, and Skyrms and Lambert, emphasize that it is not merely credences at nodes that change with evidence impact. The conceptual connections between propositions—modeled with our conditional probabilities on links—can themselves be expected to change:

[1] We use 'evidence' throughout, as opposed to 'confirmation,' as an agent- and current credence-independent measure of data impact in terms of likelihood (Brittan & Bandyopadhyay, 2019). Details are outlined in Sects. 3 and 4.

> Having reevaluated some statements we must reevaluate some others, which may be statements logically connected with the first or may be statements of logical connections themselves. (Quine 1953, p. 42)

> Here we need to pay attention not only to the probabilities of statements, but also their probabilities conditional on other statements. Conditional probabilities play a key role as rules of inference which guide belief revision, but which themselves are also modified in the process of belief revision. (Skyrms & Lambert, 1995, p. 140).

It is perhaps only now, with advanced computational capabilities applied to substantial theoretical work on Bayesian nets (Pearl, 1988, 2009), that such a vision of scientific theories and scientific change can be instantiated. We envisage scientific theories under an evidence barrage as a variety of dynamic Bayesian networks (Korb & Nicholson, 2004; Melynk, 2016; Murphy, 2002; Russell & Norvig, 2015). Here we are particularly interested in the dynamics of change within networks over time, given a continuing barrage of evidence. Different change dynamics contingent on different network structures are to be expected. But there are also constants in change dynamics that emerge as well: robust patterns of punctuated equilibrium that call to mind at least some of the phenomena of 'normal science,' 'revolutionary science,' 'paradigms' and 'paradigm shifts' emphasized in a very different philosophical tradition (Kuhn, 1969; Lakatos, 1968; Laudan, 1978).

The vagueness and multiple meanings of core concepts in that tradition, 'paradigm' among them, have long been noted to be notoriously difficult to pin down in detail (Hoyningen & Sankey, 2001; Hoyningen-Huene, 1993; Masterman, 1970; Nickles, 2003). We regard the general Kuhnian picture as a suggestive research program or working hypothesis regarding scientific history rather than anything like an established fact. What we have to offer is something very different: a structural model of scientific theory with a correlate and quantitively precise model of theoretical change. What is intriguing is that at least some of the dynamics basic to the Kuhnian picture are predictable as typical patterns of credence change within Bayesian nets under an evidence barrage. Our work can thus be seen as building and expanding on hints linking Kuhn with Bayesian networks in Salmon (1990) and Hartmann (2008).

In Sect. 2, we introduce a network conception of scientific theory. In Sect. 3, we use simple examples to outline Bayesian updating in a theoretical structure in the course of an evidence barrage. In Sect. 4, we demonstrate punctuated equilibrium in theoretical change across a variety of networks, with a further analysis of change dynamics in Sect. 5. Section 6 outlines philosophical implications. Section 7 concludes with three conjectures regarding dynamics within Bayesian models of scientific theory.

2 Theoretical structure and credence change: a Bayesian network model

Bayesian networks have been developed, interpreted, and widely applied as models of causal relations between events (Pearl, 1988, 2009; Spirtes, 2010; Spirtes et al.,

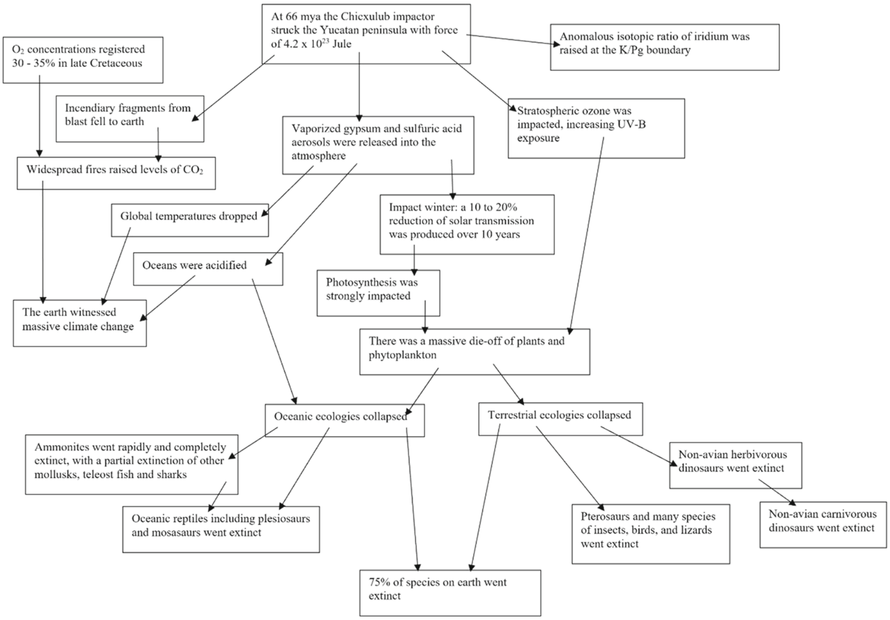

Fig. 1 A partial graph of the causal Alvarez theory regarding massive extinction at the Cretaceous-Paleogene boundary in terms of linked propositions. Sources include Jablonski and Chaloner (1994), Pope et al. (1997, 1998), Gale et al. (2001), Bardeen et al. (2017), and Henehan et al. (2019)

2000). 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 network of propositions representing causal relations between either types or tokens of events.[2] We can therefore study structural aspects of at least one type of scientific theory by applying structural lessons from Bayesian nets. We can construct a model of credence changes within a scientific theory, given a barrage of evidence over time, by studying the dynamics of change within Bayesian nets.

A partial reconstruction of the causal Alvarez theory regarding extinction at the Cretaceous-Paleogene boundary is shown in Fig. 1. Dramatic credence change percolated across this theoretical network with documentation of increased sedimentary concentrations of iridum world-wide at the Cretaceous-Paleogene boundary (Alvarez et al., 1980; Schulte et al., 2010).

There are major assumptions in many applications of Bayesian nets with an eye to causal inference: that relevant causal factors, direct or latent, are included in the representation, and that independence assumptions—that changes in one variable do not affect another—are accurate and complete. *Neither* of these requirements can be defended in the case of the representation here, offered as a partial reconstruction that captures the spirit though not the detail of a full causal theory.

[2] Pearl (2009, 2018) is particularly clear about the fluidity between interpretation of nodes as causal events or as propositions.

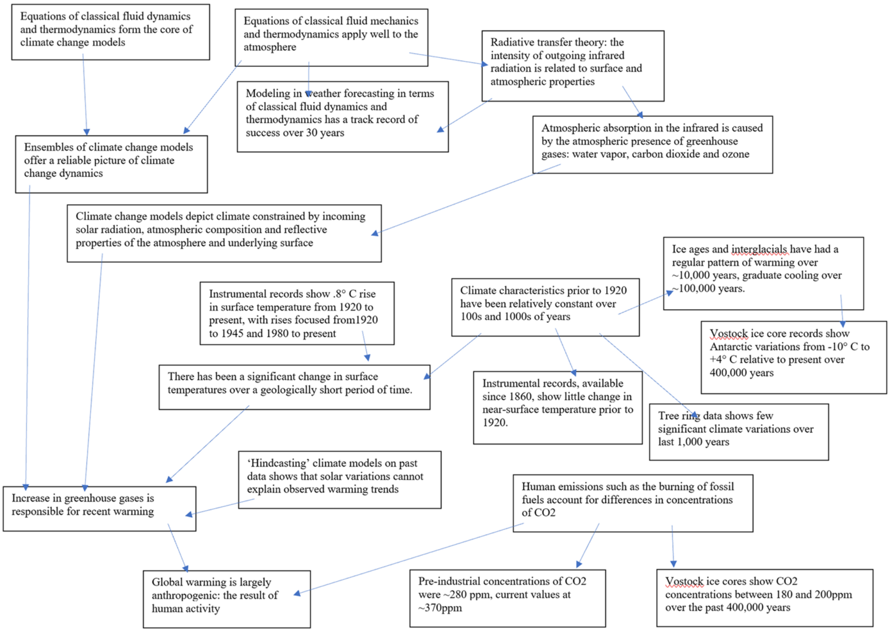

Fig. 2 A reconstruction of foundational theory regarding anthropogenic climate change. Sources include Thorpe (2005), Palmer and Stevens (2019), IPCC (2014)

Lessons from causal models can also be generalized (Climenhaga, 2019; Schaffer, 2016). 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 hypotheses 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 'x supports y as an inference' and often 'x explains y.'[3] In the other direction, evidence for more specific or applicational hypotheses can serve in Bayesian fashion to raise or lower confidence in the more general hypotheses from which they can be inferred. This direction is often intuitively clearest in terms of disconfirmation: to what extent would x be disconfirmed were its implication y disconfirmed?

A partial reconstruction of a theoretical structure of this type in the case of anthropogenic climate change is shown in Fig. 2. In this case both ice-core evidence and atmospheric data have produced important credence changes across the network (Alley, 2000; Farman et al., 1985). As in the causal case, a full Bayesian net representation would require all relevant grounding nodes and would correctly reflect all dependence and independence relations. This is certainly *not* true of the reconstruction here, a partial reconstruction that captures the spirit though not the detail of a full foundational structure.

[3] Climenhaga (forthcoming) outlines a network approach of this type informally, intended to include inference to the best explanation, enumerative induction, and analogical inference.

Classical twentieth century models of inference and confirmation were couched largely in terms of classical logic. Bayesian models offer a well-established alternative. In classical models, evidence for a hypothesis is also taken as evidence for propositions logically entailed by that hypothesis. 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 confirmation for that which logically entails that hypothesis. 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 the familiar qualitative Hempelian form.[4] 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.[5] It also offers the prospect of a detailed model of important aspects of scientific change.

Levels of credence in the elements of a scientific theory are clearly linked, and do not stand still. A single piece of evidence may confirm or disconfirm a particular element of a theoretical structure, raising or lowering credence at a node. But that credence change can also be expected to percolate through the theoretical structure, both downstream—to those elements implied by the belief in question—and upstream—to those elements in the structure that support or imply that belief. With a screening-off assumption at each step, revised credence values for the descendants of descendants—and parents of parents—can be calculated in the same fashion.[6] What we have in effect is a game of telephone, with revised credences percolating through the structure of a scientific theory as a whole.

We model a scientific theory as a directed acyclic network.[7] Nodes represent proposition-like elements which carry credence values in the open interval (0,1). The links of our model represent connections between claims within the theory. Credencechange at one point in the network produces credencechange at other 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 by Bayesian inference and conditioning on the basis of evidence at y.

But of course scientific theories are not tested once and for all with a single piece of evidence. Theoretical structures are subjected to a continuing barrage of evidence over time. We model that barrage as a series of pieces of evidence with different likelihoods at different nodes. A first piece of evidence at a node x in the network changes credence at that node in the standard Bayesian manner. That new credence percolates downstream, changing credences in the elements y of the theoretical structure

4 Both Reichenbach and the later Carnap's quantitative approaches to confirmation, on the other hand, have quite close connections with Bayesianism (Leitgeb & Carus, 2020).

5 One thing a Bayesian approach does not do, unfortunately, is solve classical paradoxes of confirmation such as Goodman's grue (Goodman, 1955). See Sober (1994) and Bandyopadhyay et al. (2016).

6 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. The importance of one-time evidential impact at different nodes within a structure, with contrasts between impact in different networks, has been developed in Grim et al. (2022).

7 Contact of our approach with either 'syntactic' or 'semantic' models of scientific theories is remote. For relation to these, and argument supporting superior virtues of a Bayesian model, see Skyrms (1984) and Hartmann (2008).

'implied' by x, using our conditional probabilities of c(y|x). Assuming screening off at y—that effects further down are only via change in y— credences in further nodes z update according to conditional probabilities c(z|y).

The rigidity condition is the assumption that conditional probabilities 'on' x such as c(y|x) or c(w|x) don't change with changes in the credence of x (Jeffrey, 1983, 1992; Shafer, 1981). But with a credence change at x, credences upstream at parents w of x change as well, and although conditional probabilities c(w|x) remain by the rigidity condition, conditional probabilities c(x|w) generally do change.

Given a first piece of evidence, then, both credences and some conditional probabilities in the network change. The new network, with those changed credence values and conditional probabilities, is then subjected to another piece of evidence, at that node or another, with changes again percolating through the structure as a whole. The result is a model of the dynamics of theoretical change under a continuing evidence barrage. Of particular interest is the fact that patterns evocative of punctuated equilibrium robustly emerge in this dynamics for a wide range of simple networks, suggestive of 'normal science' alternating with periods of 'scientific revolution' or 'paradigm shifts'.

Some major limitations should be noted, both within this model and beyond.

We work with extremely simple Bayesian networks, using point probabilities rather than probability densities, imprecise or interval-valued credences (Bovens & Hartmann, 2003; Howson & Urbach, 1989; Pearl, 1988, 2009; Pearl & MacKenzie, 2018). Our conditional probabilities update accordingly, in ways forced by the standard Bayes theorem (Sect. 3).

Independence and screening off assumptions are particularly strong in the model. Despite the graphs offered as illustration in Figs. 1 and 2, we do not assume that standard scientific theories can be formulated as Bayesian nets of the sort we use simply by treating propositions as nodes: fully formulating a theory as complicated as these in terms of complete and accurate independence and screening off conditions would be a major task. What we do assume is that theories in general will be characterized by independence and screening off conditions between propositional nodes in some given structure (Climenhaga, 2019; Henderson et al., 2010 and forthcoming); it is those structures in the abstract that we attempt to capture in a Bayesian model of the sort used here.

A limitation that is particularly relevant with regard to both the Kuhnian tradition and the attempt to model scientific change in general is the fact that the networks in our model are structurally static: neither nodes nor links are added or subtracted, though node credences can approach (though not reach) 0 or 1 and conditional probabilities can approach independence. Although we claim to model some intriguing dynamics of scientific change, we cannot claim to model them all. In that regard ours should be seen as first steps toward still richer network models of scientific change.

Beyond the models offered here, it must be admitted that there are also probabilistic network alternatives to straight Bayesian nets, well worthy of exploration (Douven & Schupbach, 2015). Our hope is that the exploration of Bayesian nets in particular may motivate further work on these alternatives as well.

3 The Bayesian mechanics of change in theoretical networks

Consider a simple theoretical structure modelled as a toy Bayesian network (Fig. 3). Given the conditional probabilities assigned (here labelled as conditional credences) the credences marked at each node are those that follow from a credence of 0.6 at the root node.

Suppose, given this structure and these credences, that we offer a piece of evidence at node b that has a likelihood of 0.25 should y be true and a likelihood of 0.75 should y be false—a piece of evidence that we will think of as having a strength of [0.75, 0.25]. On standard Bayesian conditioning, our credence at b given that piece of evidence changes from 0.58 to approximately 0.315:

$$c(b|e) = \frac{c(e|b) \times c(b)}{c(e|b) \times c(b) + c(e|{\sim}b) \times c({\sim}b)}$$

$$c(b|e) = \frac{0.25 \times 0.58}{0.25 \times 0.58 + 0.75 \times 0.42} = 0.315217$$

But of course, a change in credence at b demands a change in credences downstream as well. Given a new credence of 0.315 at b with the indicated conditional probability of c given b, and assuming the evidence e affects credence at c only by way of changed credence at b, our revised credence at c given evidence e at b changes from 0.568 to approximately 0.673:

$$c(c|e) = c(c|b) \times c(b|e) + c(c|{\sim}b) \times c({\sim}b|e)$$

by independence of c and e conditional on b

$$c(c|e) = 0.4 \times 0.315217 + 0.8 \times 0.684783 = 0.67319$$

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

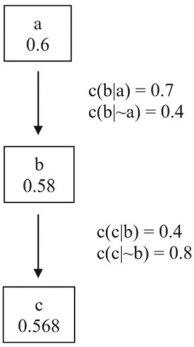

Fig. 3 A simple Bayesian network

Standardly, credence change percolates upstream as well. Using

$$c(a|b) = \frac{c(b|a) \times c(a)}{c(b)}$$

from Bayes, with our initial priors and conditional probabilities, we have

$$c(a|b) = \frac{0.7 \times 0.6}{0.58} = 0.724137931$$

Using

$$c(a|{\sim}b) = \frac{c({\sim}b|a) \times c(a)}{c({\sim}b)}$$

our initial values give us

$$c(a|{\sim}b) = \frac{0.3 \times 0.6}{0.42} = 0.428571428$$

Again using an assumption of independence, the updated credence for node a, given our new value for b on evidence e at b is

$$c(a|e) = c(a|b) \times c(b|e) + c(a|{\sim}b) \times c({\sim}b|e)$$

by independence of a and e conditional on b

$$c(a|e) = 0.724137931 \times 0.315217 + 0.428571428 \times 0.684783 = 0.521739$$

Thus, credence at parent node a changes from 0.6 to approximately 0.522.

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.[8]

With the assumption of a piece of evidence of strength [0.75, 0.25] at node b, our credences at a, b, and c have changed from <0.6, 0.58, 0.568> to <0.522, 0.315, 0.673>. But this is not all that changes in our network with the impact of evidence. Bayesian networks are often used for one-shot calculations: 'given set conditional probabilities and evidence that the grass is wet, what must the probability be that it has rained?' Scientific theories, on the other hand, are subjected to a continuing barrage of evidence over time. Under a continuing evidence barrage, not only credences at individual nodes in a Bayesian network of the form used here but conditional probabilities between those nodes will change. These changes are forced formally by demands of network consistency and Bayes' theorem. In spirit, they are entirely in line with the general picture of scientific change outlined in Quine, Skyrms and Lambert.

[8] On network propagation both 'bottom-up' and 'top-down' see Lauritzen and Spiegelhalter (1988) and Normand and Tritchler (1992).

By the rigidity condition, probabilities conditional on b and ~b do not change: c(a|b) and c(a|~b) in this example, as well as c(c|b) and c(c|~b). But the probabilities c(b|a) and c(b|~a) on which b is conditional, do change; and indeed by standard Bayesian principles they must. Our new value for a is 0.315; our new value for b is 0.673. But were the conditional probabilities c(b|a) and c(b|~a) to retain their original values, our value for b would be not 0.673 but

$$\begin{aligned} c\ (b|e) &= c(b|a) \times\ c\ (a|e) +\ c(b|{\sim}a) \times\ c\ ({\sim}a|e) \\ &= 0.7 \times 0.521 + 0.4 \times 0.479 = 0.5563 \end{aligned}$$

Were there no change in probabilities conditional on b, in other words, we'd be saddled with an inconsistent network.

What Bayes requires is that we replace our initial conditional probabilities c(b|a) and c(b|~ a) with posterior conditional probabilities which we will term $c'(b|a)$ and $c'(b|{\sim}\ a)$.[9] Here we use c(a|b) and c(a|~b), calculated as above, which by rigidity do *not* change. By Bayes,

$$\begin{aligned} c'(b|a) &= c(a|b) \times\ c(b|e)/c(a|e) = 0.437499566 \\ c'(b|{\sim}a) &= c({\sim}a|b) \times\ c(b|e)/c({\sim}a|e) = 0.181817906 \end{aligned}$$

Use of these posterior conditional probabilities gives us a consistent network with our value of 0.315 for b|e:

$$\begin{aligned} c(b|e) &=\ c'(b|a) \times\ c(a|e) +\ c(b|{\sim}a) \times\ c({\sim}a|e) \\ &= 0.4375 \times 0.521 + 0.1818 \times 0.479 = 0.315 \end{aligned}$$

Evidence of strength [0.75, 0.25] at node b has therefore not merely changed credences at other nodes, but in terms of conditional probabilities has changed the network as a whole. This updating for conditional probabilities is effectively incorporated in standard programming for Bayesian networks such as aGrum and pyAgrum, which we've employed for simulations throughout (Ducamp et al., 2020; Gonzales et al., 2017).

Incorporating posterior credences and posterior conditional credences, what we might think of as our posterior network appears in Fig. 4.

As a scientific theory, of course, such a network is subject to further evidence as well. We suppose evidence of strength [0.35, 0.65] at node c. Given that evidence and this network, using the same Bayesian calculations as above, our network is transformed again. That new network appears on the extreme right of Fig. 5, which tracks network changes through our two evidence 'hits.'

Scientific theories are not tested with a single piece of evidence, once and for all. They are subject to a continuing incoming barrage of evidence, with credence in individual elements of the theory and in the connection between them changing over time. Our attempt is to model what the change dynamics of theoretical networks look

9 Here we are particularly grateful to Jim Joyce.

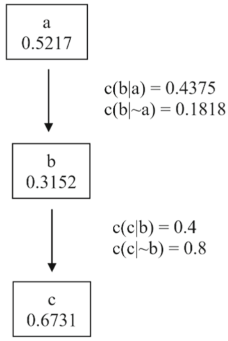

Fig. 4 Transformation of network in this figure given evidence of strength [0.75, 0.25] at node b, with values rounded off to the fourth decimal place

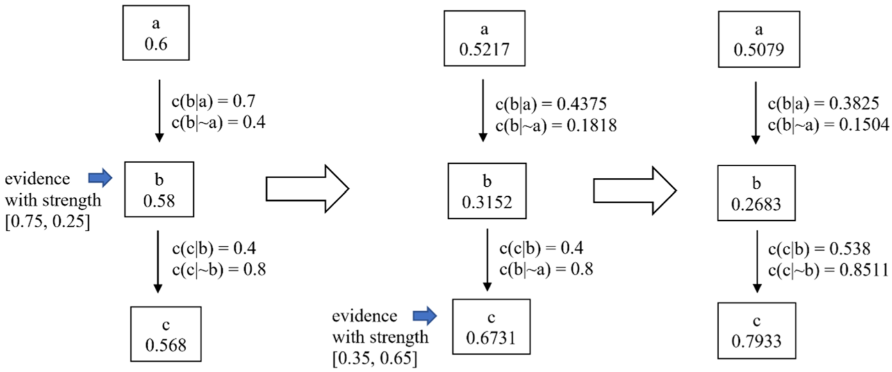

Fig. 5 Transformation of network in given successive evidence inputs, values rounded the fourth decimal place

like under a continuing barrage of evidence—a longitudinal study within a Bayesian model very much in the spirit of Quine, Skyrms and Lambert.

In the calculation above we have tracked credence changes node by node. As a metric of credence change across a network as a whole, we can use Brier divergence (Brier, 1950). 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.[10] Formally:

[10] There are of course alternative measures. Brier divergence instantiates a difference measure between posterior and prior credences, |c(x|e) − c(x)|, as a measure of impact. An alternative would be to build on a ratio measure c(x|e)/c(x). Although alternative measures would of course change the absolute values of change, in neither this example nor in general would that appear to make any difference in the general patterns of network change under an evidence barrage. 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) offers a very useful overview of the latter.

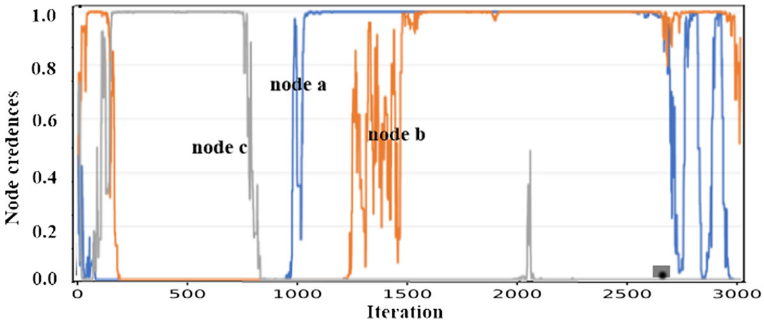

Fig. 6 For the network used in Figs. 3 through 5, node credences under a random evidence barrage. 3000 iterations shown

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

where p and q are probability measures. In the case of the first evidence impact illustrated in Fig. 5 above, the network moves from credences of <0.6, 0.58, 0.568> to <0.5217, 0.3152, 0.6739> , with rounded values used for simplicity. The Brier divergence from step one to step two is therefore approximately 0.0292. From step two to step three credences start at the previous network transformation values of <0.5217, 0.3152, 0.6739> and shift to <0.5079, 0.2683, 0.7933> , giving us a Brier divergence of approximately 0.0056.

4 Network transformations under an evidence barrage

What does a pattern of network transformations look like under a continuing barrage of evidence over time? Here we continue to hit the network outlined above with such a barrage, choosing a random node at each iteration and a piece of evidence of strength [1 − x, x] for a random real value x in the interval (0.1, 0.9). Were we to draw from a full (0, 1) interval we would include effects that might be due merely to occasional pieces of overwhelming evidence, since it is the Bayes factor in terms of the ratio of the two elements of the vector that dictates evidence strength (Kass & Raffery 1995). By clipping the ends of a (0, 1) interval to (0.1, 0.9) we avoid that potential distraction.[11]

At each impact we update both credences and conditional probabilities as outlined above, representing the new network that will be subjected to the next piece of evidence. A first way of envisaging network changes is in terms of changing credences at each node. For one series of 3000 random evidence hits, the pattern of node credences in the network outlined above is shown in Fig. 6.

The corresponding pattern in terms of Brier divergence across the network as a whole is shown in Fig. 7.

[11] Here we are grateful for the suggestion of an anonymous referee.

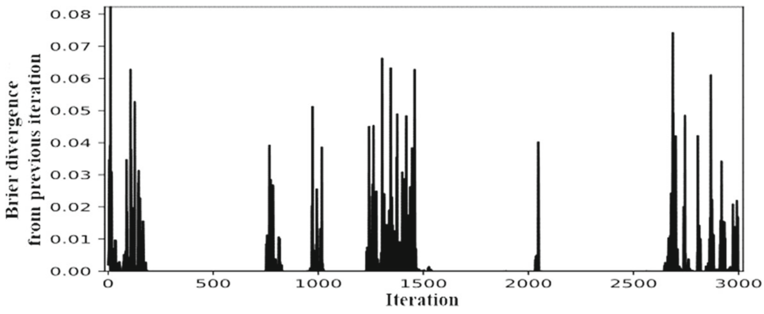

Fig. 7 For the network used in Figs. 3 through 5, Brier divergence under the random evidence barrage used in Fig. 6. 3000 iterations shown

Credence changes and Brier divergence for the same network under another random evidence barrage is shown in Figs. 8 and 9.

Our focus in what follows is the clear qualitative features of these typical patterns, an analysis of their dynamics, and what we take to be the philosophical significance

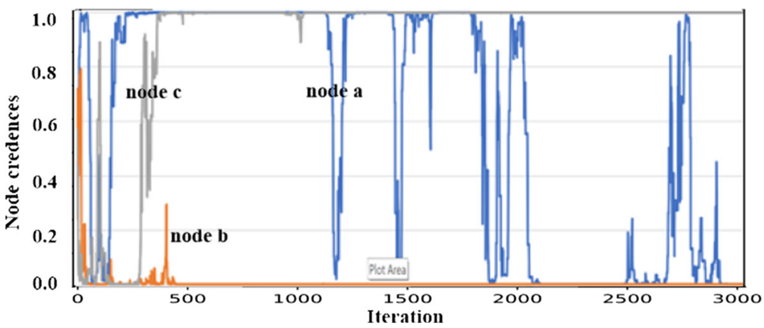

Fig. 8 For the network used in Figs. 3 through 5, node credences under a second random evidence barrage. 3000 iterations shown

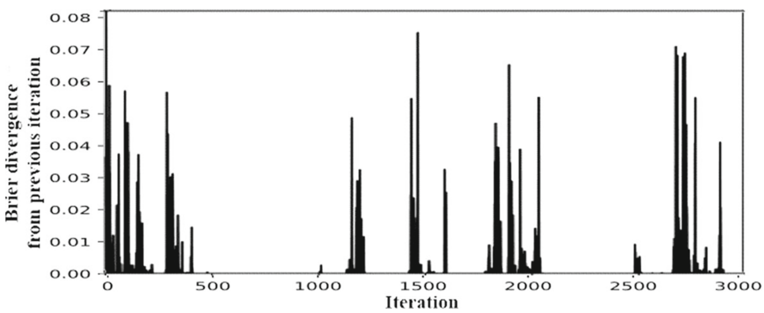

Fig. 9 For the network used in Figs. 3 through 5, Brier divergence under the random evidence barrage used in Fig. 8. 3000 iterations shown

of these with regard to theoretical change. Of particular note is the fact that these graphs show (a) distinctive periods of intense activity in terms of divergence between iterations under random evidence impact, together with (b) clear periods of flat calm, with very little divergence between iterations under evidence impact. Although evident in terms of credence change graphed node by node, patterns are particularly clear when viewed in terms of Brier divergence across the network as a whole. In Figs. 10 and 11 we extend iterations from 3000 to 30,000 in each case, demonstrating that with different intervals the qualitative pattern of punctuated equilibrium continues.

We also offer a slightly more complex theoretical network—the all-too-familiar water sprinkler example of a Bayesian net, here cast as a simple causal theory (Fig. 12).

The initial credences shown for each proposition in Fig. 12 are those that follow from a credence of 0.5 at the root node using the conditional probabilities shown. These values are standard for this familiar example except that we have avoided conditional values of full strength, 0 or 1. In this example we again subject the theoretical network to a barrage of successive evidence of strength [1 − x, x] for a random real value x in the clipped interval (0.1, 0.9). Both credences and conditional probabilities change as outlined. Brier divergences from iteration to iteration for this theoretical structure under a continuing evidence barrage are shown in Fig. 13.

In the water sprinkler example, as in the simple linear network, patterns of change under an evidence barrage show distinct periods of (a) intense activity in terms of divergence from iteration to iteration together with (b) periods of calm: a pattern of punctuated equilibrium (Eldredge & Gould, 1972, 2007; Gould & Eldredge, 1977).

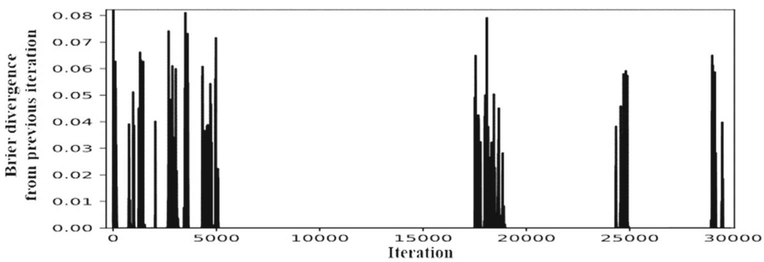

Fig. 10 Brier divergence under the random evidence barrage used in Figs. 6 and 7. 30,000 iterations shown

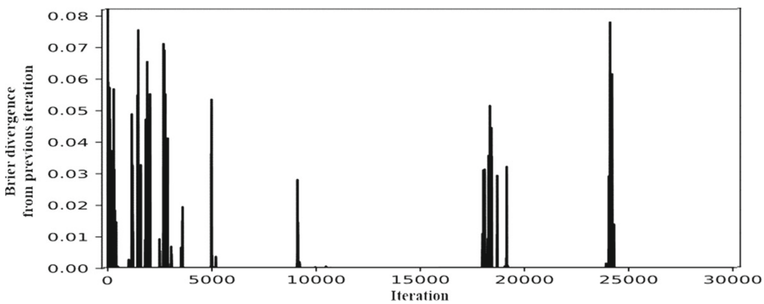

Fig. 11 Brier divergence under the random evidence barrage used in Figs. 8 and 9. 30,000 iterations shown

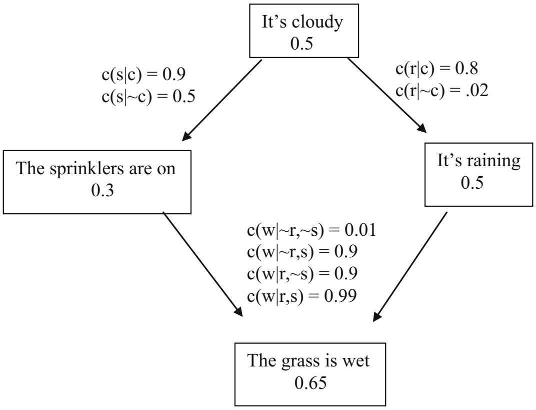

Fig. 12 The ubiquitous Bayesian water sprinkler example, here as a simple causal theory

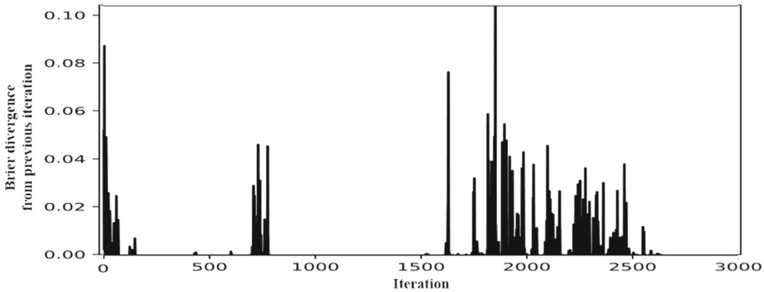

Fig. 13 Brier divergence between iterations in the water sprinkler network in Fig. 10 under a random evidence barrage

The evidence barrages we've used to this point consist of pieces of evidence of random strength impacting randomly selected nodes. Evidence impacting randomly selected nodes, however, might be thought to be counter-intuitive in light of our interpretation of networks as theoretical structures. A theoretical structure might be thought of as 'impinging on experience at its edges,' in Quine's phrase, with impacts primarily at the 'periphery.' In Fig. 14 we use the same series of random evidence strengths, but hit our nodes with a graduated probability toward the periphery. In this case the leaf node 'grass is wet' is impacted eight times out of every thirteen, the middle nodes each two times out of thirteen, and the root node only once every thirteen times. As is to be expected, the pattern of change in Fig. 13 differs from Fig. 14. But here too a qualitative pattern of periods of relative stasis interrupted by periods of intense divergence is clear.

In the next section we offer a simple explanation for this punctuated equilibrium dynamic, followed by an outline of the philosophical lessons suggested by such a model.

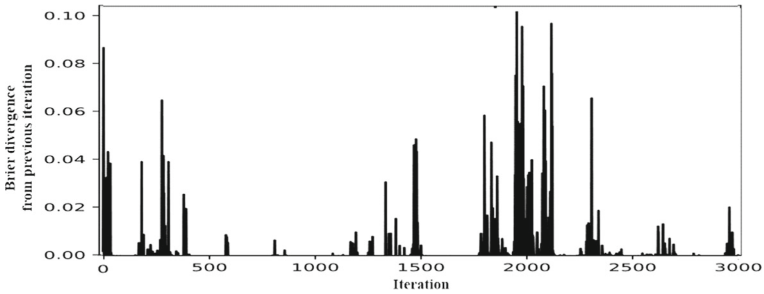

Fig. 14 Brier divergence between iterations in the water sprinkler network with a random evidence barrage targeted differentially toward the leaf node

5 Understanding punctuated equilibrium in Bayesian nets

Why does a random evidence barrage produce a punctuated equilibrium pattern of credence change in Bayesian nets? It is clear from the figures above that the specific pattern of change dynamics—the specific pattern of periods at which the network is quiescent and at which it is volatile—depends on the specifics of the random evidence series to which a network is subjected. But regardless of the specifics of the random series of nodes and evidence values used as the evidence barrage, patterns of punctuated equilibrium appear robustly in the change dynamics. Why?

A potential answer is not far to seek. On standard Bayesian principles, a node with credence close to an extreme of either 0 or 1 will be more resistant to change in the face of evidence than will a node with credence farther from these extremes. By the same token, a network comprised of nodes close to either 0 or 1 will be relatively resistant to change when presented with evidence.[12] We thus offer the explanatory hypothesis that the quasi-equilibrium periods in the history of a Bayesian net under evidential barrage are those points at which its credences approach extremes of 0 or 1; the punctuated periods of volatility are those points at which its credences have been nudged from those extremes.

Figure 15 offers confirmation for such a hypothesis. For each iteration in Fig. 7, we graph the Brier divergence of the network at that point from a 'constant distribution' network in which each node takes a credence of either 0 or 1, whichever is closest to that node's current value. Higher values in Fig. 15 thus represent networks farther from node values of 0 and 1; lower values represent networks closer to node values of 0 and 1. Comparison shows that stable periods in Fig. 7 correspond to periods in which node distances from 0 or 1 are vanishingly small; periods of volatility correlate with periods in which at average node distance is further from those polar values. An overlay of Figs. 7 and 15, emphasizing the point, appears as Fig. 16.

Figure 17 offers a similar confirmation in the water sprinkler case in Fig. 13, showing a clear correspondence between (a) Brier divergence from a 'constant distribution' of

[12] Though we emphasize the role of extreme credences here, a related point will hold for conditional probabilities: conditional probabilities such that p(a|b) = p(a|~ b) will preserve stasis, whereas those which vary from that equality will favor change.

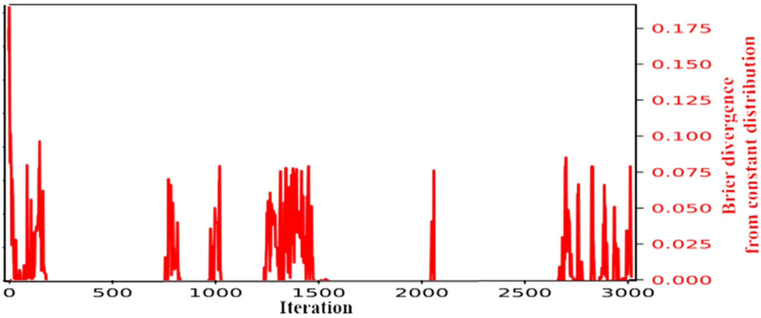

Fig. 15 For comparison with Fig. 7: Brier divergence at each iteration of **a** the credence values of the network shown in Fig. 6 from **b** a constant distribution in which node credence are replaced with whichever is closer, 0 or 1

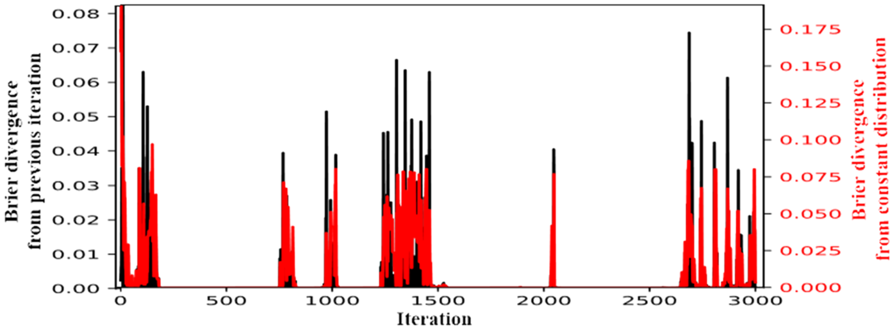

Fig. 16 An overlay of Brier divergence from previous iteration (black) as in Fig. 7, with Brier divergence from a constant distribution (red) as in Fig. 15

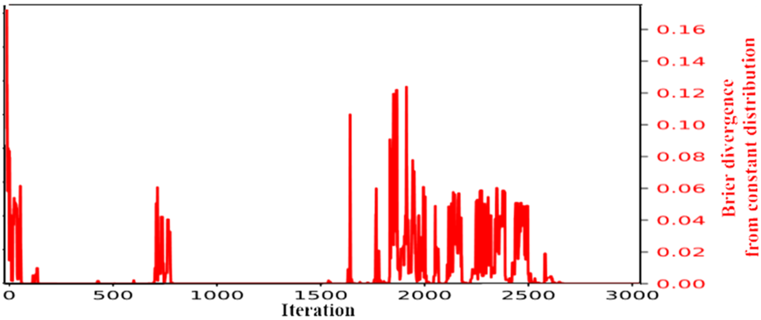

Fig. 17 For comparison with Fig. 13: Brier divergence at each iteration of **a** the credence values of the water sprinkler network shown in Fig. 9 from **b** a constant distribution in which node credence are replaced with whichever is closer, 0 or 1

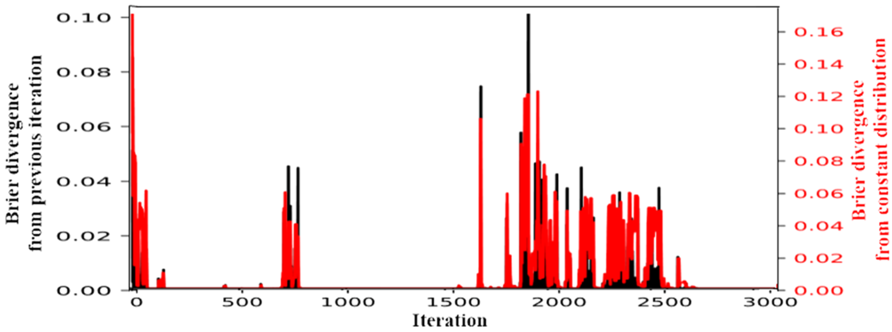

Fig. 18 An overlay of Brier divergence from previous iteration (black) as in Fig. 13, with Brier divergence from a constant distribution (red) as in Fig. 17

node values of 0 and 1 and (b) volatility in the dynamics of credence change in the network. An overlay of Figs. 13 and 17, emphasizing the point, appears as Fig. 18. A similar correspondence between (a) and (b) holds for the other examples, and in every case that we have explored in the course of investigation.[13]

It should be noted that a network that settles down to node values close to 0 and 1 at different points of stasis doesn't always settle down to the *same* network node values. In the dynamics of our three-node network shown in Figs. 6 and 7, for example, node values <a, b, c> at iteration 500 approach <0, 0, 1>. At iteration 900, they approach <0, 0, 0> instead. At iteration 1600 they approach <1, 1, 0>. Figure 19 uses three dimensions, one for each node, to graph the credence values of the three nodes in the course of the evidence barrage. In the supplementary materials we offer a movie of credence value change through 30,000 iterations of this barrage.

There may be an analog far stronger than mere metaphor between mechanisms of punctuated equilibrium in biological evolution and the patterns of change we trace in Bayesian networks under an evidence barrage. With diffusion equations modeling probability distributions of mean phenotypes with genetic drift on a landscape with several adaptive peaks, Lande (1985) builds on work by Simpson (1944, 1953) and suggestions in Wright (1977), showing a pattern of long periods of relative stasis interrupted with short periods of change between peaks:

> The time for the final transition between adaptive peaks is expected to be orders of magnitude shorter than the time the population initially spends drifting around the original adaptive peak… (Lande, 1985 p. 7644)

Barton and Charlesworth (1984) show a similar pattern with a formula taken from chemical physics.

These evolutionary models lack both the network structure and the specific Bayesian dynamics of the present model, and are of course applied to biological evolution rather than scientific theory change. What they have in common with the current model are

[13] In cases in which we exempt nodes entirely from evidential impact, periods of stasis may occur when those exempted nodes are not near 0 or 1. In these cases we have a higher but constant line in periods of stasis in the equivalents of Figs. 11 and 12 simply because of a constant non-extreme value in the exempted node.

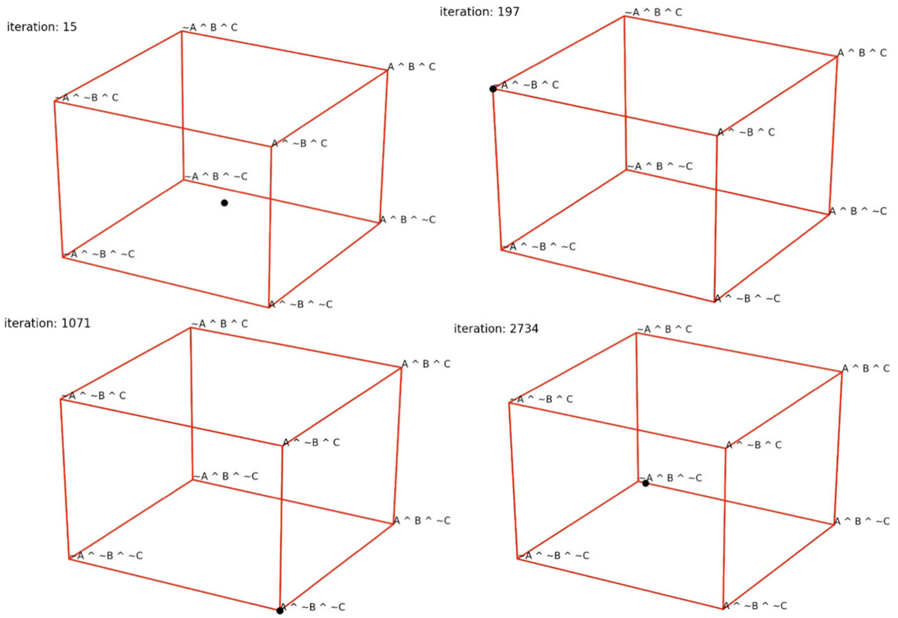

Fig. 19 A three-dimensional representation of credence values for the three nodes of the network used in Figs. 3 through 5 in the course of the evidence barrage shown in Figs. 6 and 7, 15 and 16. A complete animation is available in the supplementary materials

periods of 'stickiness,' at adaptive peaks in the biological case and extreme priors in the case of Bayesian networks. It is that stickiness that shows up as periods of stasis, with quick shifts between those periods on the basis of either dynamic. There are also tantalizing hints that the network structure of the current model may echo *co*-evolutionary structures that may be of importance in the dynamics and rates of evolutionary change (Taper & Case, 1992).[14]

6 Normal science, paradigm shifts, and the punctuated equilibrium of scientific change

We suggest that a Bayesian network model of the structure of scientific theory, plausible in its own terms, offers a new understanding of at least some of the phenomena of scientific change that have been the foci of a very different, purely qualitative, and historically oriented tradition in the philosophy of science. At least some of the phenomena at the core of the Kuhnian tradition are predictable without elaborate historical context, predictable in the typical dynamics of scientific theory change captured simply as Bayesian nets under a barrage of incoming evidence. One might think that a (bordered) random evidence barrage of the sort employed in the current model would produce an equally random-looking pattern of credence changes in the network: noise

[14] We are grateful to an anonymous referee for bringing these models to our attention.

in and noise out. What our results indicate is that this is not the case. In response to even a bordered random input, the dynamics of credence change in Bayesian nets result in a pattern of punctuated equilibrium.

We should emphasize that the model seems to capture *some* Kuhnian phenomena. The vagueness and multiple meanings of 'paradigm' in the Kuhnian tradition have long been remarked upon, together of course with companion concepts of 'paradigm shift,' 'normal science,' and 'revolutionary science' (Masterman, 1970). Emphasis is sometimes put on paradigms as sets of scientific habits, experimental traditions, or complexes of puzzle-solving techniques. In this sense Kuhn analyzes tests for 'goodness of air' in Priestley and Lavoisier as paradigms, for example (Kuhn, 1969, 59–60). More centrally, however, the tradition glosses 'paradigms' as sets of received beliefs (Kuhn 4–5), views of nature (2–3), successful theories (17–18), or the scientific achievements of major bodies of theory: "Aristotle's *Physica*, Ptolemy's *Almagest*, Newton's *Principia* and *Optics*, Franklin's *Electricity*, Lavoisier's *Chemistry* and Lyell's *Geology*" (10). It is something like paradigms in this latter sense that our model best captures. At least some of the basic dynamics predicted in the Kuhnian model regarding paradigms in this sense, dynamics commonly emphasized in historical examples, and dynamics that continue to motivate researchers in the Kuhnian tradition, appear robustly and unforced in a simple Bayesian model of scientific theory and credence change.

If scientific theories or cognitive systems in general *do* have something like the structure of Bayesian nets, even a random barrage of evidence can be expected to produce something far from random: a punctuated equilibrium characterized by periods of relatively little credence change—the characteristic of 'normal science'—interrupted by periods of credence volatility—characteristic of 'scientific revolutions' or 'paradigm shifts.'[15]

7 Three conjectures

We conclude with two conjectures regarding application of the current model to other networks and a third conjecture concerning the robustness of punctuated equilibrium across model variations.

Conjecture 1 In general, networks that are more complex in terms of number of nodes will show longer periods of both stasis and volatility on the current model.

This seems plausible on the grounds that a longer series of random evidence 'hits' will generally be required to drive a network with more nodes to a point at which all node values approach 0 and 1. By the same token, a longer evidence barrage may be required to drive a network that is close to fixation into an alternative. This is not

[15] With regard to punctuated equilibrium, Gould himself characterizes Kuhn as "The most influential punctuationalist theory in twentieth century scholarship" (Gould, 2007, p. 276) and credits his reading of Kuhn as a major impetus behind the evolutionary theory (Gould, 1989, 2007). It has not escaped our notice that the dynamics tracked for Bayesian networks offers an intriguing explanation for punctuated equilibrium in evolutionary change as well: an ecological network explanation as an alternative to the traditional explanation in terms of allotropic speciation (Eldredge & Gould, 1972; Mayr, 1954; Simpson, 1944, 1953).

of course to deny that there will be cases where change in a single node can cause a cascade across even a large and complex network.

To the extent that conjecture 1 holds, on the model for scientific theories offered here, more complex and elaborated theories will show longer periods of little change in the manner of 'normal science,' but also longer periods of punctuated 'scientific revolution' in the form of volatile credence change.

Conjecture 2 In general, more sparsely connected networks—those with fewer links between nodes—will be more resistant to change on the current model, showing longer periods of both stasis and volatility.[16]

This seems plausible because the impact of change at a node will percolate to those nodes to which it is connected, 'decaying' in impact to nodes beyond nodes. The more connected a network, the more the direct impact of change at any one node will tend to be.

On the model for scientific theories offered here, if both conjectures 1 and 2 hold up it will be theoretical structures elaborated in terms of many propositions but with fewer implicational links between them that will show longer periods of both 'normal' and 'revolutionary' science.

Conjecture 3 The qualitative patterns of punctuated equilibrium documented here will prove robust across major model variations.

As noted initially, the model offered here involves some major limitations. We have worked throughout with simple Bayesian nets employing point probabilities, but conjecture that very similar results will hold for more sophisticated models using interval-valued credences. Any scientific network can be expected to involve independence and screening off conditions, but those used in our sample networks are particularly strong. We conjecture that similar results will appear with different patterns of independence and screening off, though precisely how dynamics vary with those aspects of structure is a large question calling for further work. We have also worked throughout with random evidence within a wide but bounded range, far wider than can be expected in general for evidential variance and a 'shaking hand' of inaccuracy. Our informal explorations of far narrower ranges of randomness continue to show punctuated equilibrium, though with smaller swings in Brier divergence and more widely separated periods of stasis.

Our conjecture that qualitative results will prove generally robust is grounded in the explanation for observed dynamics offered in Section 5. The punctuations of punctuated equilibrium are due to accumulated effects of variations in incoming evidence. Equilibria form as Bayesian credences are pushed nearer the poles and thus more resistant to change. Although the details can be expected to change, the basic dynamics of punctuated equilibrium can be expected to hold across Bayesian models in general.

There are other limitations in the current model which this does not address. There are for example probabilistic network alternatives to the Bayesian nets explored here (Douven & Schupbach, 2015). We consider these well worthy of exploration. It remains

[16] As noted by a referee, the extreme case of a network with no connections would be immune to contagious changes.

an open but intriguing question to what extent the patterns of theoretical change documented here will characterize alternative plausible models as well.

A limitation that we've noted as particularly relevant with regard to both the Kuhnian tradition and the attempt to model scientific change is the fact that the networks within our model are structurally static: though both credences and conditional probabilities can change radically, neither nodes nor links are added or subtracted. Better models are needed in order to capture that aspect of scientific change as well.

8 Conclusion

An attractive vision of scientific theories as networks of propositions—as webs of belief—comes to us from Quine (1951, 1953), with some first steps toward formalization in Skyrms and Lambert (1995). Contemporary network theory offers the prospect of a deeper and more detailed understanding of how evidence sensitivity, 'falsifiability,' and propensity to change play out across different forms of theory envisaged as different propositional networks. An invitation to focus on the *dynamics* of scientific change comes to us from the Kuhnian tradition, with an appealing but vague vision of normal science punctuated by periodic scientific revolutions (Kuhn, 1969; Lakatos, 1968; Laudan, 1978). Bayesian and other forms of updating dynamics on networks offer the prospect of more precise models of scientific stability and change, across a range of metrics, and at least in the abstract.

Ours is intended as a first step in the exploration of the dynamics of scientific theory change using Bayesian network models. Using Brier divergence from iteration to iteration as a metric, change dynamics within such a model show a distinct pattern of punctuated equilibrium under random and periphery-targeted evidential barrages. Our suggestion is that these precise and quantitative phenomena, derivable from network structure alone, capture at least some of the phenomena long noted qualitatively in the Kuhnian tradition of normal science and paradigm shifts. The conjectures we've offered, if true, would extend both that formal analysis and its philosophical implications.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11229-022-03720-z.

Acknowledgements We are grateful to Brian Skyrms, Jim Joyce, Mark Newman, Gordon Belot, and Josh Hunt for help and encouragement, and to three anonymous referees for careful review, well-placed criticism, and constructive suggestions.

Funding n/a.

Data availability Data availability of data and coding from authors on request.

Declarations

Conflict of interest The authors declared that they have no conflict of interest.

References

Alley, R. B. (2000). Ice-core evidence of abrupt climate changes. *PNAS, 97*(4), 1331–1334.

Alvarez, L. W., Alvarez, W., Asaro, F., & Michel, H. V. (1980). Extraterrestrial cause for the Cretaceous-Tertiary extinction. *Science, 208*(4448), 1095–1108.

Bandyopadhyay, P. S., Brittan Jr., G., & Taper, M. L. (2016). The paradoxes of confirmation. In *Belief, evidence, and uncertainty: Problems of epistemic inference*. Springer briefs in philosophy (pp. 125–141). Springer.

Bardeen, C. G., Garcia, R. R., Toon, O. B., & Conley, A. J. (2017). On transient climate change at the Cretaceous-Paleogene boundary due to atmospheric soot injections. *Proceedings of the National Academy of Sciences of the United States of America, 114*(36), 201708980.

Barton, N. H., & Charlesworth, B. (1984). Genetic revolutions, founder effects, and speciation. *Annual Review of Ecology, Evolution and Systematics, 15*, 133–164.

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

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review, 78*(1), 1–3.

Brittan, G., & Bandyopadhyay, P. S. (2019). Ecology, evidence, and objectivity: In search of a bias-free methodology. *Frontiers in Ecology Ad Evolution, 7*, 399. https://doi.org/10.3389/fevo.2019.00399

Climenhaga, N. (forthcoming). Evidence and inductive inference. In M. Lasonen-Aarnko & C. Littlejohn (Eds.), *The Routledge handbook of the philosophy of evidence*. Routledge.

Climenhaga, N. (2019). The structure of epistemic probabilities. *Philosophical Studies*. https://doi.org/10.1007/s11098-019-01367-0

Crupi, V., Tentori, K., & Gonzalez, M. (2007). On Bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science, 74*(2), 229–252.

Douven, I., & Schupbach, J. N. (2015). Probabilistic alternatives to Bayesianism: The case of explanationism. *Frontiers in Psychology, 6*, 459. https://doi.org/10.3389/fpsyg.2015.00459

Ducamp, G., Gonzales, C., & Wuillemin, P.-H. (2020). aGrUM/pyAgrum: A toolbox to build models and algorithms for probabilistic graphical models in Python. In *Proceedings of the 10th international conference on probabilistic graphical models, PMR 138*. http://proceedings.mlr.press/v138/ducamp20a.html

Eldredge, N., & Gould, S. J. (1972). "Punctuated equilibria: An alternative to phyletic gradualism. In T. J. M. Schopf (Ed.), *Models in paleobiology* (pp. 82–115). Freeman Cooper.

Farman, J. C., Gardiner, B. G., & Shanklin, J. D. (1985). "Large losses of total ozone on Antarctica reveal seasonal ClO_x/NO_x interaction. *Nature, 315*(6016), 207–210.

Gale, J., Rachmilevitch, S., Reuveni, J., & Volokita, M. (2001). The high oxygen atmosphere toward the end-Cretaceous: A possible contributing factor to the K/T boundary extinctions and to the emergence of C_4 species. *Journal of Experimental Botany, 52*(357), 801–809.

Gonzales, C., Torti, L., & Wuillemin, P.-H. (2017). aGrUM: A graphical universal model framework. In *International conference on industrial, engineering, and other applications of intelligent systems*. https://doi.org/10.1007/978-3-319-60045-1_20

Goodman, N. (1955). *Fact, fiction, and forecast*. Harvard University Press.

Gould, S. J. (1989). Punctuated equilibrium in fact and theory. *Journal of Social and Biological Structures, 12*, 117–136.

Gould, S. J. (2007). *Punctuated equilibrium*. Belknap Press.

Gould, S. J., & Eldredge, N. (1977). Punctuated equilibria: The tempo and model of evolution reconsidered. *Paleobiology, 3*(2), 115–151.
Grim, P., Seidl, F., McNamara, C., Rago, H. E., Astor, I. N., Diaso, C., & Ryner, P. (2022). Scientific theories as Bayesian nets: Structure and evidence sensitivity. *Philosophy of Science, 89*(1), 42–69.
Hartmann, S. (2008). Modeling in philosophy of science. In M. Frauciger & W. K. Essler (Eds.), *Representation, evidence, and justification: Themes from Suppes*. Lauener library of analytic philosophy (Vol. 1, pp. 95–122). ontos Verlag.
Henderson, L., Goodman, N. D., Tenenbaum, J. B., & Woodward, J. F. (2010). The structure and dynamics of scientific theories: A hierarchical Bayesian perspective. *Philosophy of Science, 77*, 172–200.
Henehan, M. J., Ridgwell, A., Thomas, E., Zhang, S., Alegret, L., Schmidt, D. N., Rae, J. W. B., Witts, J. D., Landman, N. H., Greene, S. E., Huber, B. T., Super, J. R., Plnavsky, N. J., & Hull, P. M. (2019). Rapid ocean acidification and protracted Earth system recover following the end-Cretaceous Chicxulub impact. *Proceedings of the National Academy of Sciences of the United States of America, 116*(45), 22500–22504. https://doi.org/10.1073/pnas.1905989116
Howson, C., & Urbach, P. (1989). *Scientific reasoning: The Bayesian approach*. Open Court.
Hoyningen, P., & Sankey, H. (Eds.). (2001). *Incommensurability and related matters*. Kluwer.
Hoyningen-Huene, P., 1993. *Reconstructing scientific revolutions: Thomas S. Kuhn's philosophy of science*. University of Chicago Press.
IPCC. (2014). Summary for policymakers. In C. B. Field, V. R. Barros, D. J. Dokken, K. J. Mach, M. D. Mastrandrea, T. E. Bilir, M. Chatterjee, K. L. Ebi, Y. O Estrada, R. C Genova, B. Girma, E. S. Kissel, A. N. Levy, S. MacCracken, P. R. Mastrandrea, & L. L White (Eds.), *Climate change 2014: Impacts, adaptation and vulnerability. Part A: Global and sectoral aspects. Contribution of working group II to the fifth assessment report of the Intergovernmental Panel on Climate Change* (pp. 1–32). Cambridge University Press.
Jablonski, D., & Chaloner, W. G. (1994). Extinctions in the fossil record [and discussion]. *Philosophical Transactions: Biological Sciences, 344*(1307), 11–17.
Jeffrey, R. C. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.
Jeffrey, R. C. (1992). *Probability and the art of judgment*. Cambridge University Press.
Kass, R. E., & Raffery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90*(430), 773–795.
Korb, K. B., & Nicholson, A. E. (2004). *Bayesian artificial intelligence*. Chapman and Hall/CRC.
Kuhn, T. S. (1969). *The structure of scientific revolutions* (4th ed.). University of Chicago Press.
Lakatos, I. (1968). Criticism and the methodology of scientific research programmes. *Proceedings of the Aristotelian Society, 69*, 149–186.
Lande, R. (1985). Expected time for random genetic drift of a population between stable phenotypic states. *Proceedings of the National Academy of Sciences of the United States of America, 82*, 7641–7645.
Laudan, L. (1978). *Progress and its problems: Towards a theory of scientific growth*. University of California Press.
Lauritzen, S. L., & Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society Series B, 50*, 157–224.
Leitgeb, H., & Carus, A. (2020, forthcoming). Rudolf Carnap. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2020 ed.). https://plato.stanford.edu/archives/fall2020/entries/carnap/
Masterman, M. (1970). The nature of a paradigm. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge.* Cambridge University Press.
Melynk, I. (2016). Dynamic Bayesian networks: Estimation, inference, and applications. Dissertation, University of Minnesota.
Murphy, K. P. (2002). Dynamic Bayesian networks. A book chapter based on his dissertation, U. C. Berkeley. Retrieved March 5, 2021, from https://www.cs.ubc.ca/~murphyk/Papers/dbnchapter.pdf
Nickles, T. (Ed.). (2003). *Thomas Kuhn*. Cambridge University Press.
Normand, S.-L., & Tritchler, D. (1992). Parameter updating in a Bayes network. *Journal of the American Statistical Association, 98*(420), 1109–1115.
Palmer, T., & Stevens, B. (2019). The scientific challenge of understanding and estimating climate change. *Proceedings of the National Academy of Sciences of the United States of America, 116*(49), 24390–24395.
Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Morgan Kaufman.
Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.
Pearl, J., & MacKenzie, D. (2018). *The book of why*. Basic Books.

Pope, K. O., Baines, K. H., Ocampo, A. C., & Ivanov, B. A. (1997). Energy, volatile production, and climatic effects of the Chicxulub Cretaceous/Tertiary impact. *Journal of Geophysical Research: Planets, 102*(E9), 21645–21664.

Pope, K. O., D'Hondt, S. L., & Marshall, C. R. (1998). Meteorite impact and the mass extinction of species at the Cretaceous/Tertiary boundary. *Proceedings of the National Academy of Sciences of the United States of America, 85*(19), 11028–11029.

Quine, W. V. O. (1951, 1953). Two dogmas of empircisim. *The Philosophical Review, 60*(1), 20–43. Reprinted in *From a logical point of view*. Harvard University Press, 1953.

Russell, S. J., & Norvig, P. (2015). *Artificial intelligence: A modern approach* (3rd ed.). Pearson.

Salmon, W. (1990). Rationality and objectivity in science or Tom Kuhn meets Tom Bayes. In C. W. Savage (Ed.), *Scientific theories* (pp. 175–204). University of Minnesota Press.

Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies, 173*(1), 49–100.

Schulte, P., et al. (2010). The Chicxulub asteroid impact and mass extinction at the Cretaceous-Paleogene boundary. *Science, 327*(5970), 1214–1218.

Shafer, G. (1981). Constructive probability. *Synthese, 48*(1), 1–60.

Simpson, G. G. (1944). *Tempo and mode in evolution*. Columbia University Press.

Simpson, G. G. (1953). *The major features of evolution*. Columbia University Press.

Skyrms, B. (1984). *Pragmatics and empiricism*. Yale University Press.

Skyrms, B., & Lambert, K. (1995). The middle ground: Resiliency and laws in the web of belief. In F. Weinert (Ed.), *Laws of nature: Essays on the philosophical, scientific and historical dimensions.* De Gruyter.

Sober, E. (1994). No model, no inference: A Bayesian primer on the grue problem. In D. Stalker (Ed.), *Grue! The new riddle of induction* (pp. 225–240). Open Court.

Spirtes, P. (2010). Introduction to causal reasoning. *Journal of Machine Learning Research, 11*, 1643–1662.

Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction and search* (2nd ed.). MIT Press.

Taper, M. L., & Case, T. J. (1992). Models of character displacement and the theoretical robustness of taxon cycles. *Evolution, 46*(2), 317–333.

Thorpe, A. J. (2005). Climate change prediction: A challenging scientific problem. Institute of Physics. https://www.iop.org/publications/iop/archive/page_52088.html#gref

Wright, S. (1977). *Experimental results and evolutionary deductions* (Vol. 3). University of Chicago Press.