Philosophy of Science, Network Theory and Conceptual Change:
Paradigm Shifts as Information Cascades

Patrick Grim, Joshua Kavner, Lloyd Shatkin, & Manjari Trivedi

Forthcoming in Euell Elliot and L. Douglas Kiel, eds., *Complex Systems in the Social and Behavioral Sciences: Theory, Method, and Application*, Ann Arbor: Univ. of Michigan Press

Abstract

Philosophers have long tried to understand scientific change in terms of a dynamics of revision within 'theoretical frameworks,' 'disciplinary matrices,' 'scientific paradigms' or 'conceptual schemes.' No-one, however, has made clear precisely how one might model such a conceptual scheme, nor what form change dynamics within such a structure could be expected to take. In this paper we take some first steps in applying network theory to the issue, modeling conceptual schemes as simple networks and the dynamics of change as cascades on those networks. The results allow a new understanding of two traditional approaches—Popper and Kuhn—as well as introducing the intriguing prospect of viewing scientific change using the metaphor of self-organizing criticality.

Introduction

The attempt to understand science, its dynamics, and how it changes might call for any of various levels of analysis—and must ultimately include them all. A psychologist might approach the issue with an eye to creativity and conformity. An economist might approach the topic in terms of incentives for research, for innovation or exploitation of existing resources. A sociologist might think of the task as primarily a social study, concentrating on research communities, structures of journal communication, and academic procedures for advancement and funding.

Philosophers have typically thought of particular areas of scientific research as characterized by 'theoretical frameworks,' disciplinary matrices,' 'scientific paradigms,' or 'conceptual schemes' at a particular time, with scientific change to be understood as changes in those conceptual structures. It is presumed that such schemes exist psychologically in some individual head, that they are shared and changed through the social dynamics of science, and that change may follow a form of both individual and social incentive. The philosopher's level of analysis, however, is the 'theoretical framework,' 'disciplinary matrix,' scientific paradigm,' or 'conceptual scheme' itself, envisaged as something like an abstract object. The philosopher's presumption is that at least major aspects of scientific change will be understandable as broadly logical and boundedly rational changes in the scheme itself, though those changes play out in epistemic economics, through psychological mechanisms and in a social dynamics. Under pressure of new evidence, theoretical frameworks and scientific paradigms can be expected to change. The philosopher's goal, at that level of analysis, is to better understand how.

The work we offer here triangulates from familiar work by the two most influential and familiar philosophers of science of the 20th century: Karl Popper and Thomas S. Kuhn (Popper 1959, 1963; Kuhn 1962, 1969, 1970). If the influence of an academic discipline is measured by the extent to which its concepts are incorporated within the wider culture, 20th century philosophy of science can boast of only two such successes. The influence of logical positivism, logical empiricism, Carnap, Hempel, Goodman, Quine and Bayesian epistemology has been largely *intra*-discipline (Carnap 1966, 1969; Hempel 1966; Goodman 1955; Quine 1951, 1953; Bovens & Hartmann 2003). Only two conceptual configurations have crossed disciplines to enter wider scientific discourse and the discourse of society at large. Both of those influential philosophical achievements are theories of scientific change.

Popper speaks of hypothesis and observation as always presupposing "a frame of reference: a frame of expectations: a frame of theories…a theoretical framework" (Popper 1963, p. 62). Kuhn is speaking at the same level of analysis when he refers to a 'disciplinary matrix' … a set of "shared beliefs, values, instruments, and techniques" (Kuhn 1969, p. 174). His other term for a 'disciplinary matrix,' of course, is a 'scientific paradigm.' Popper and Kuhn clearly agree that in order to understand the dynamics of science we need to understand change within theoretical frameworks or disciplinary matrices: the structure and change of scientific paradigms.

The basic idea of paradigms has a long and distinguished history. The idea that science and cognition in general operate not in terms of isolated elements but a system is as old as Plato's *logos* in the Theatetus, Aristotle's *scientia* in the *Posterior Analytics*, the Medieval *summa*, and explicit attention to system in Descartes, Malebranche, Spinoza, Newton, and Leibniz. The concept continues beyond Kuhn in Imre Lakatos's research programmes (Lakatos 1968), Larry Laudan's research traditions (Laudan 1978), Willard van Orman Quine's 'webs of belief' (Quine & Ullian 1978) and widespread contemporary references to 'conceptual schemes,' 'conceptual frameworks,' and 'scientific worldviews'

Despite this long philosophical tradition, thinking in terms of paradigms, theoretical frameworks, and conceptual schemes, however, no-one in the philosophical tradition has attempted to model a conceptual scheme or track its dynamics. How precisely might one model a paradigm? How might one track, even theoretically, the dynamics of paradigm shift?

Our attempt here is to take some first steps by putting the tools of complex systems and network theory in particular to use in philosophy of science. The science of complexity, we argue, carries important philosophical lessons regarding the complexity of science.

I. Scientific Paradigms: A First Model

How might one model a paradigm? As first step, consider a set of claims—the claims of a scientific theory, or of linked scientific theories within a discipline. Those claims together compose the theoretical framework or scientific paradigm operative at a particular time. But of course those parts of a paradigm don't float independently: the claims of such a framework are linked by broadly logical connections. Some claims follow from others. Some are read as evidential instantiations of others. Some are read as theoretical generalizations. It is a set of claims bound by connections of mutual support that constitute a scientific paradigm. It is the intuition that we can model scientific paradigms as networks that guides our work throughout.

We start with a model of mutual support within elements of a paradigm using the simplest possible model: an undirected graph. With even that simple picture, however, we can model two clear aspects of scientific dynamics. Science grows. And science changes.

Consider a set of nodes with no connections. Those are disparate observations, perhaps, hypotheses and conjectures, but have not yet been integrated into a systematic body of theory.

As a science develops, those claims become further integrated. They start to form a larger whole—something that begins to deserve the name of science—the conceptual system that constitutes a paradigm. That process of progressive integration, long observed in historical episodes of science, can be modelled by progressively adding links within our model. As a discipline matures its integration increases: the network grows new links (Fig. 1).

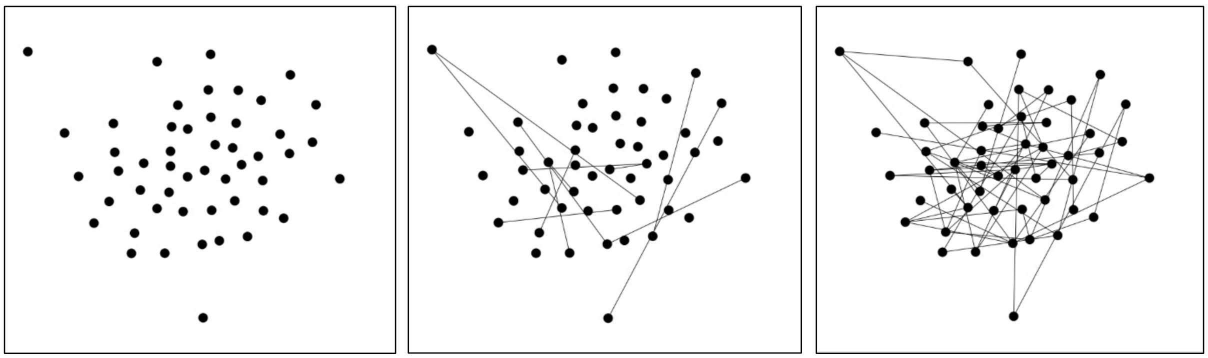

Fig. 1 Modeling progressive integration within a network

But of course there are other aspects of scientific change. Theories not only rise but fall. Paradigms form, but are also liable to collapse: the phenomenon of paradigm shift.

II. Modeling Popperian Falsification

How does science change? On the Popperian picture, science is a matter of conjectures in search of refutations. Targeting advocates of Freud and Marx, Popper argues that finding apparent confirmations or verifications of a set of beliefs is all too easy. With Einstein and the Eddington expeditions as a favored example, Popper argues that the mark of genuine science is not safety in vagueness but risk in precision. The fundamental mechanism of scientific changes is crucial experimentation in which conceptual systems are falsified (Popper 1959, 1963).

In a Popperian model, we let all nodes of our modeled conceptual system start with an 'established' value: 1, for convenience. But falsifications happen. For modeling purposes, we suppose falsifications occur at random, somewhere in the conceptual structure. Because a Popperian 'theoretical structure' is a structure of linked nodes, those falsifications may well spread. If an observational consequence of a given claim is falsified, it is not merely that claim but any claim that entails it—any higher elements of theory—that are falsified as well. Falsification can be expected to spread through the conceptual structure (Fig. 2).

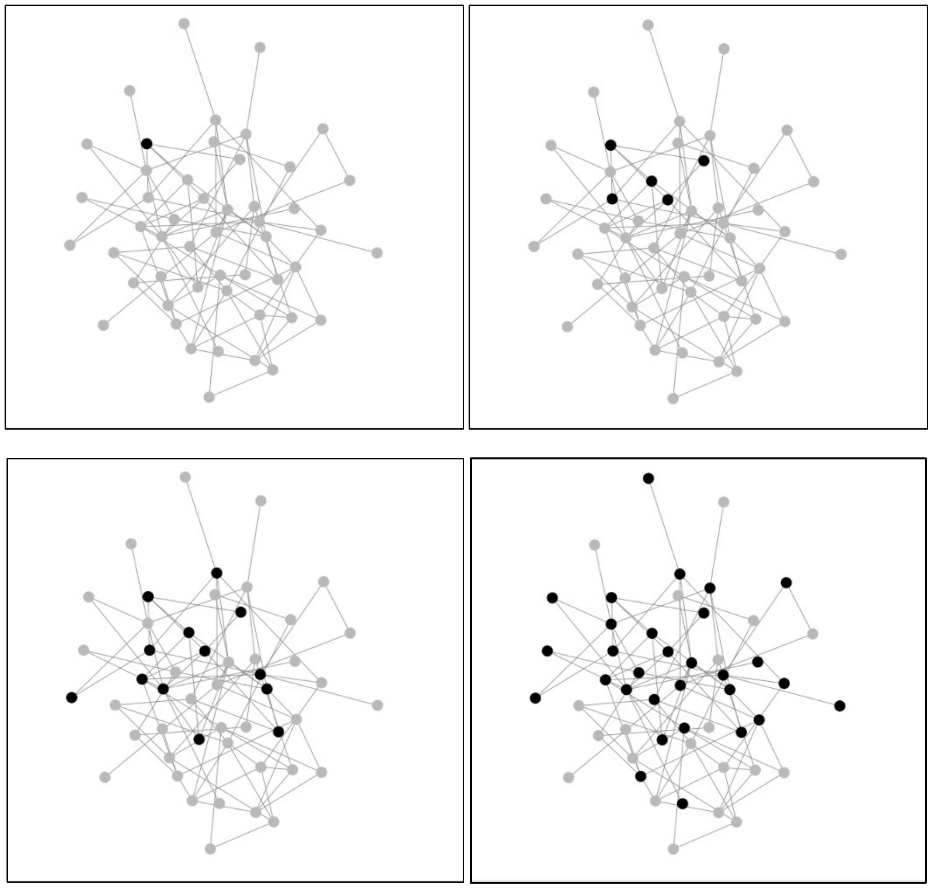

Fig. 2 Spreading falsification within a Popperian model

What can we expect scientific change to look like on the Popperian picture? That is a question within the philosophy of science. A model allows us to make it more specific. With a theoretical framework or conceptual system modeled as an undirected network, how much of that network can we expected to be affected by a random falsification? How often and how large can we expect falsification cascades to be?

The basic question is one in philosophy of science. But with a model in hand, it is complex systems and specifically network theory that offers us an answer. The sizes of falsifiability cascades in a Popperian network will depend on the connectivity of the network—its characteristic degree or average number of links per node—and will depend on characteristic degree in very interesting ways.

We start with data from simulation, turning to the analytic background for explanation.

We start with a Popperian network of 50 nodes in which the average degree of our conceptual nodes is less than 1. We consider 1000 random networks, each with characteristic degree—average node degree—of .5. Within each of those networks we drop a falsification at a random point in the structure. We then track the sizes of 'cascades': the number of points out of 50 that must fall by Popperian falsification spreading from that initial spot (Fig. 3).

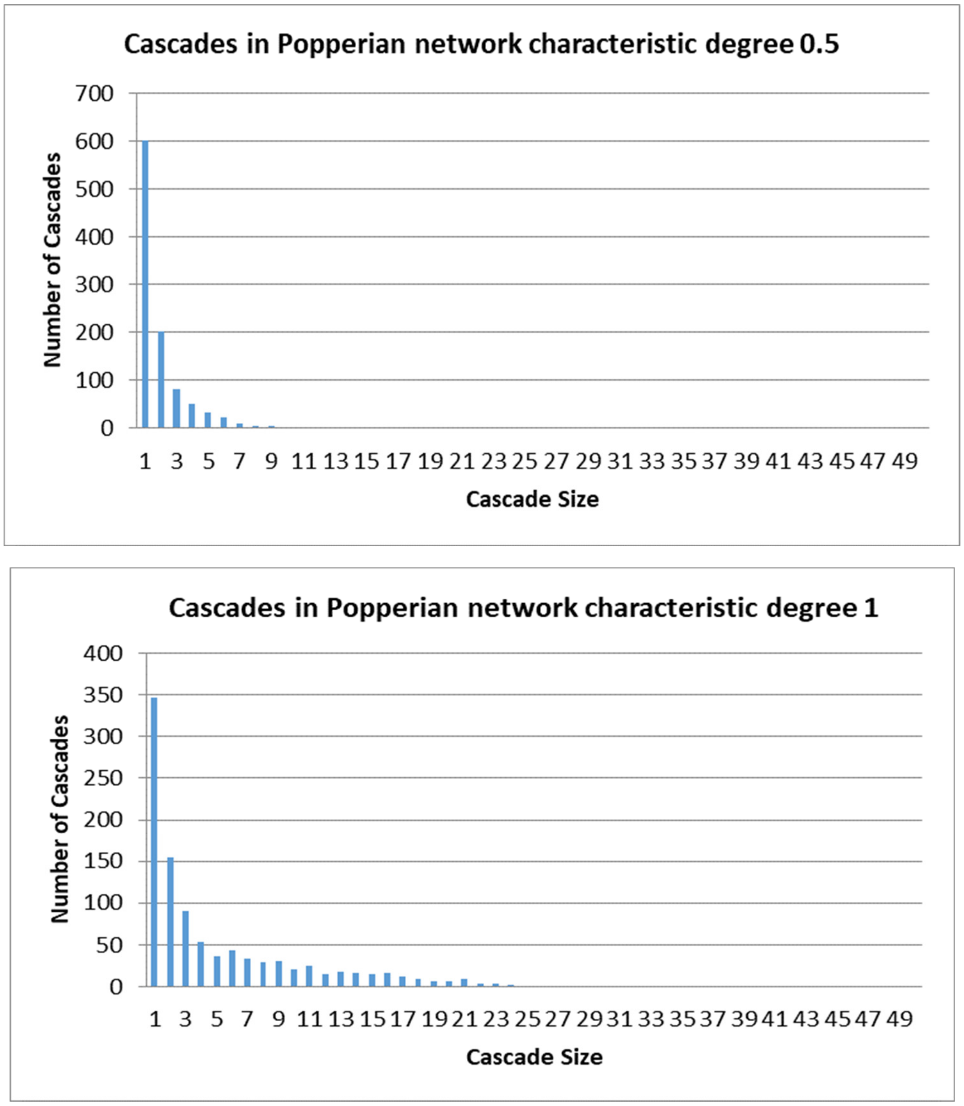

Fig. 3 Distribution of cascade sizes in Poppereian networks of characteristic degree 0.5 and 1. Note change in scale on y-axis.

Falsification cascades on networks with a characteristic degree of 1 or less show something like a power law distribution. That pattern changes dramatically for Popperian networks have greater integration, marked by higher characteristic degrees. At a characteristic degree of 1.5 far fewer drops are isolated, with a distribution that clearly becomes bimodal. At characteristic degree 2 the pattern is even more noticeable. Here the cascade sizes that are most frequent are cascades of 41 or 42 nodes, with the bimodal distribution clearly dominated by large cascades at the right. At that degree of integration and above, the characteristic result of a single falsification is a widespread cascade throughout a major part of the network (Fig. 4).

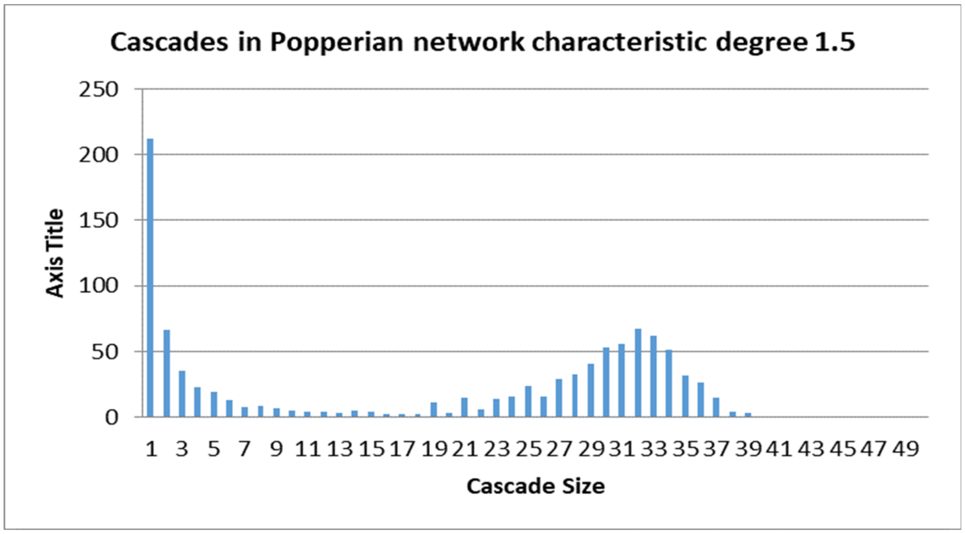

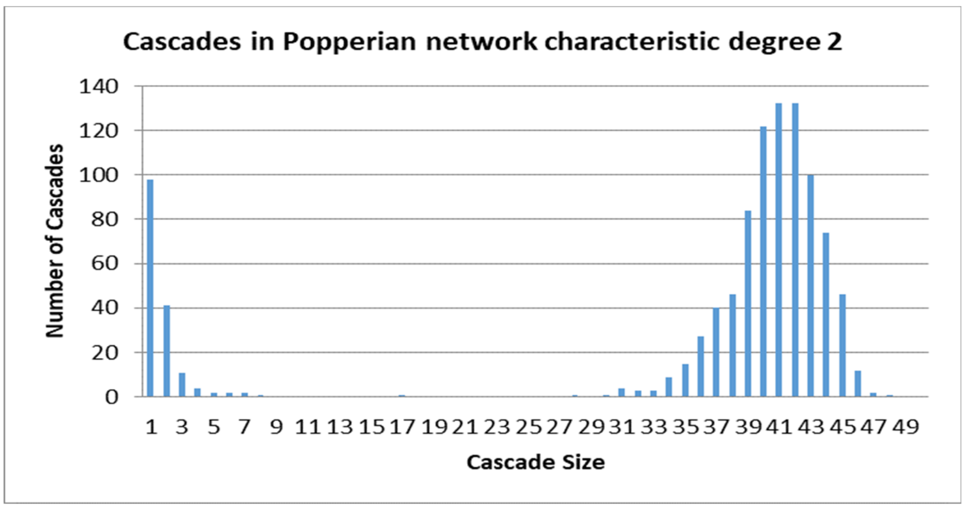

Fig. 4 Distribution of cascade sizes in Popperian networks of characteristic degree 1.5 and 2. Note change in scale on y-axis.

Although the specific application of these results lies within philosophy of science, the formal basis lies within classic network theory. The formal basis for these results is the phenomenon of giant components, outlined in Erdös and Rényi's classic work on random graphs (Erdös and Rényi 1959, Newman 2010). Figure 5 shows the familiar graph of a phase transition to a giant component at a characteristic degree 1, with rapid increase in proportion of the network as characteristic degree increases.

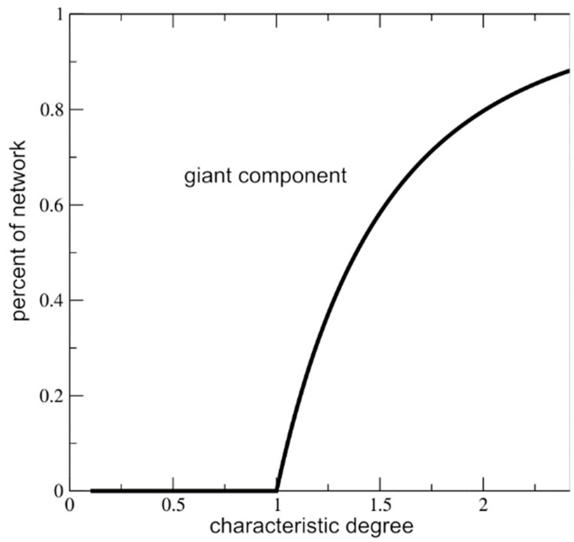

Fig. 5 Formation of a giant component as percent of a network at different degrees within an undirected random network. After Erdös and Rényi 1959.

The crucial fact in application to philosophy of science is that a falsification in a Popperian network will affect all of the nodes to which it is connected. If a 'giant component' occupies a certain proportion of the network, the probability that a single falsification will cascade through that proportion of the network will be the probability that an arbitrary node is within the giant component. That probability, of course, will correspond to the proportion of the network occupied by the giant component.

In order to gauge the potential size of falsification cascades on Popperian conceptual networks, in other words, all we really need to know is the characteristic degree of the network in question. Results correspond directly to the graphs of cascade sizes I've shown you. Another aspect of this analysis, to which we will return in section VI, is that the proportion of the network occupied by a giant component is scale free, independent of network size n (Erdos & Renyi, Newman, Watts?)

III. Modeling Kuhnian Dynamics

Despite the fact that Kuhn and Popper are often portrayed as antagonists, Kuhn writes quite explicitly of all that they have in common (Kuhn 1970). But Kuhn emphasizes that scientific change is rarely if ever a matter of decisive falsification on the basis of a single crucial experiment. A single anomaly—an unexplained phenomenon or apparent piece of counter-evidence—is never fatal. In 1827 Robert Brown noticed through a microscope that grains of pollens dance on the surface of water. But Brownian motion remained simply an unexplained curiosity until incorporated into theories of molecular motion (Perrin 2005). Copernican theory predicts stellar parallax: two closely separated stars ought to appear closer to each other at some times than at other times. But from Copernicus's time well into the nineteenth century no stellar parallax was observed—an anomaly that was shrugged off using an auxiliary hypothesis regarding the limits of available telescopes (Curd 1982).

Established paradigms resist change, maintained by explaining away apparent counter-evidence, impugning the expertise of critics, building ad hoc supplementary hypotheses. But anomalies build up. No single anomaly is fatal. But anomalies can accumulate in such a way as to weaken confidence in one part of a theory, or one aspect of a paradigm. That in turn can weaken confidence in another part of the theory or paradigm. The build-up of anomalies across different areas of a paradigm can lead to crisis, signaling an imminent paradigm shift.

We can model the Kuhnian picture of scientific change by replacing simple values of 1 and 0 for our nodes with thresholds: the number of accumulated anomalies at which a specific component of the network will be abandoned. The failure of one portion of a paradigm—a node or set of nodes—will pass the pressure of anomaly to other portions. Often, as Kuhn notes, an anomaly may remain localized, without affecting the structure as a whole. But in some contexts there will be a cascade of anomalies across the structure: the mark of a paradigm shift.

We assign random node thresholds between 1 and 5, and then begin dropping 'anomalies' at random. In Kuhnian fashion, anomalies accumulate, so we retain the same network throughout a consecutive sequence of drops. When a node reaches its threshold, it drops anomalies to each of the other nodes to which it is connected. The result may be simply a small number of anomalies distributed locally. But under some conditions node after node may reach anomaly threshold, producing a cascade of anomaly-forced change across the network.

It should be emphasized that a Kuhnian model differs in major ways from the Popperian. There we dropped a single falsification on 1000 different networks. Here we drop 1000 cumulative anomalies on a single network, though we then average results over 100 such networks.

There are a number of variations possible in a Kuhnian model. Here we offer two. In a first case, when a node reaches its full threshold value it passes a single anomaly to the nodes to which it connected (one anomaly passed). In this form of the model we measure a cascade in terms of the number of nodes that reach full threshold from an anomaly drop on a random node (a full tipping cascade). We zero-out a node's anomalies when it has 'discharged' an anomaly down the line. It becomes a new node, as it were, in a new paradigm.

Paradigm shift cascades on a Kuhnian network, like falsification cascades on a Popperian network, will depend on the integration of the network. Results for a random network of degree 1 are shown in Figure 6. Most single anomalies will fail to tip even that node on which they drop. The highest number of drops will result in a single tip, with fewer that result in a cascade that tips two nodes, fewer still that tip three. The result is a power law distribution at low node degree very much like that for Popperian networks.

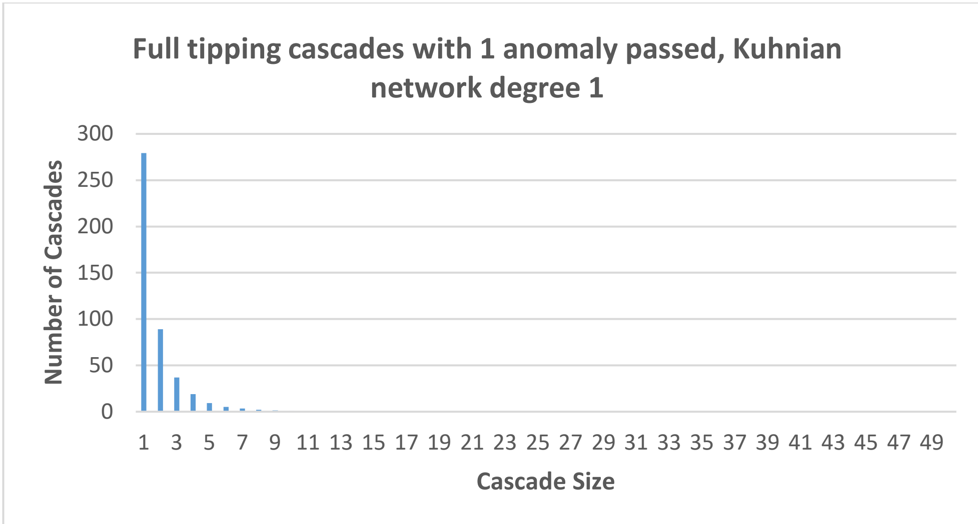

Fig. 6 Cascades of fully tipped nodes in a Kuhnian network of degree 1, with 1 anomaly passed to neighboring nodes at threshold.

Results with the same mechanism for Kuhnian networks of higher degree are shown in Figure 7. Despite the fact that the Kuhnian model uses a very different mechanism than the Popperian—1000 successive anomalies dropped on single networks, rather than single falsifications dropped on 1000 different networks—the qualitative character of cascade distributions is very similar. At low degree the pattern shows a power law. At higher degree a clear bimodal pattern emerges, with global cascades across a large portion of the network.

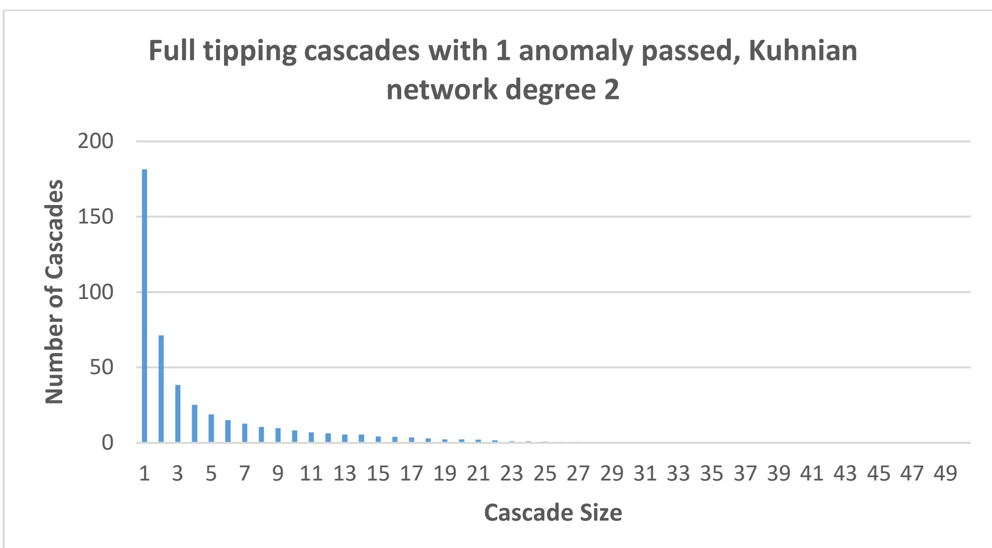

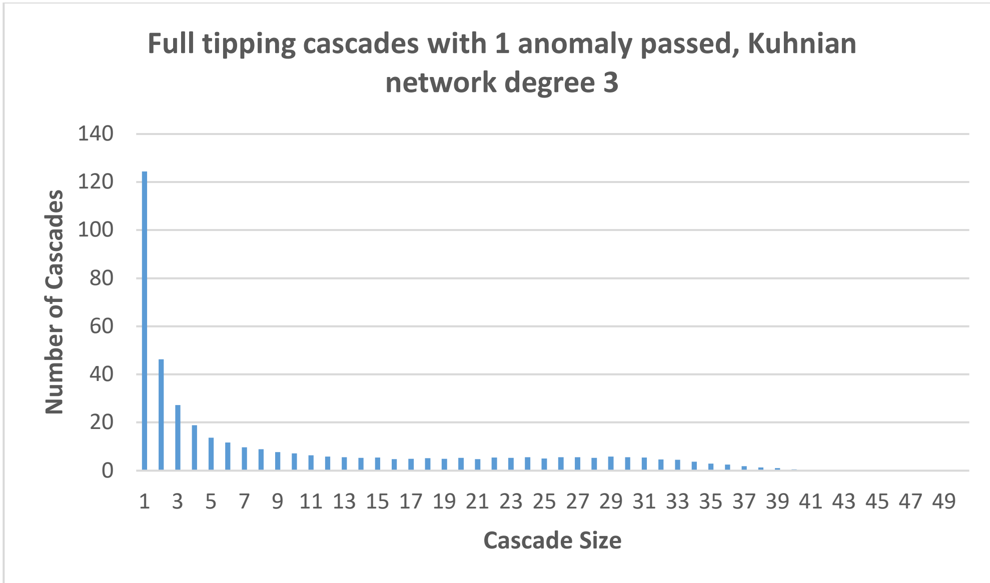

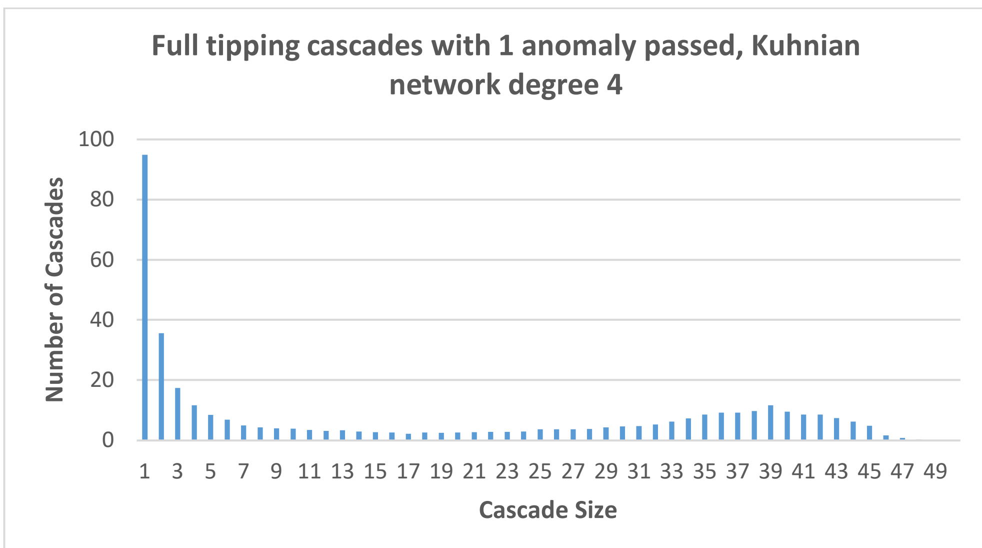

Fig. 7 Cascades of fully tipped nodes in a Kuhnian network of degree 1, with 1 anomaly passed to neighboring nodes at threshold.

Patterns of global cascade in threshold networks roughly analogous to the Kuhnian model remain a topic of continuing research (Watts 2002, Chang & Lyuu 2009, Hurd, Gleeson & Melnick 2017). As Watts notes,

> Global cascades in social and economic systems, as well as cascading failures in engineering networks, display two striking qualitative features: they occur rarely, but by definition are large when they do. This general observation, however, presents an empirical mystery. Both power-law and bimodal distributions of cascades would satisfy the claim of infrequent, large eents, but these distributions are otherwise quite different, and might require quite different explanations. (Watts 2002, 5771)

It is clear both here and in that continuing research that the clue to large cascades in threshold networks is not simply the formation of a giant component, as in the case of Popperian networks, but the formation of a giant *vulnerable* component: a connected set of nodes all of which are a small component shy of threshold and thus ready to tip. A clue to the qualitative similarity between Popperian and Kuhnian networks is the fact that the Popperian can be seen as a limiting case of the Kuhnian: a Kuhnian model in which node thresholds are set uniformly to 1.

As noted, there are a range of variations possible in a Kuhnian model. In a second variation we have nodes pass their full threshold values on saturation, rather than a single anomaly. Here we measure a cascade as the number of nodes beyond the drop spot the contents of which are changed, whether or not those nodes reach full threshold. With that different concept of cascades an intriguingly different pattern appears (Fig. 8).

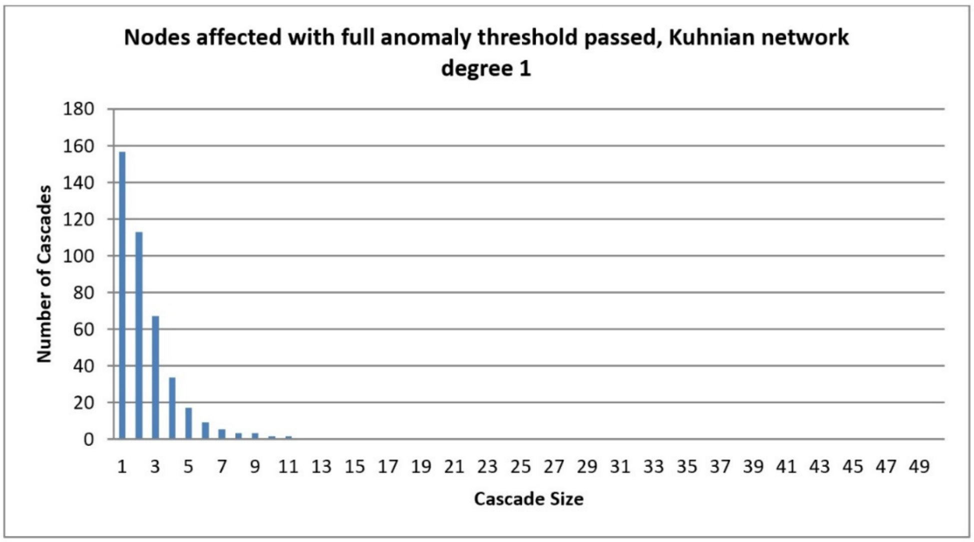

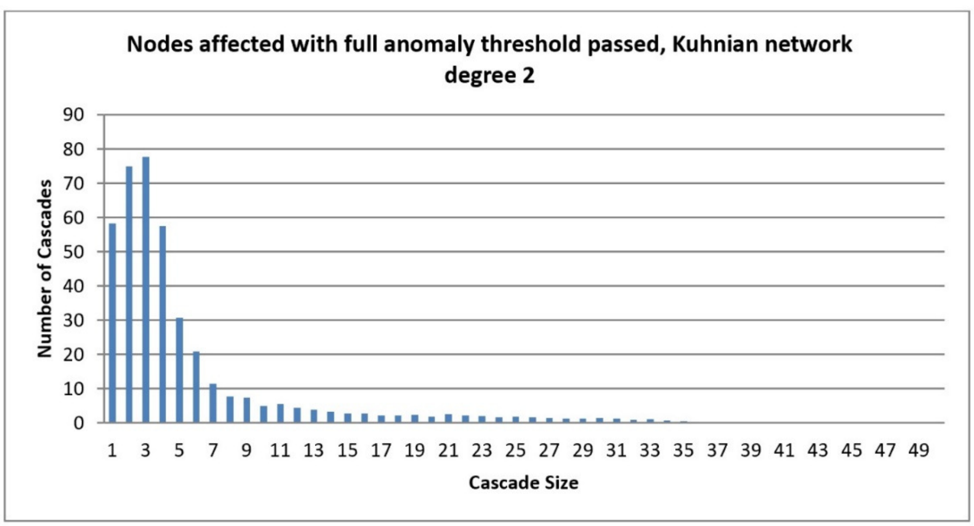

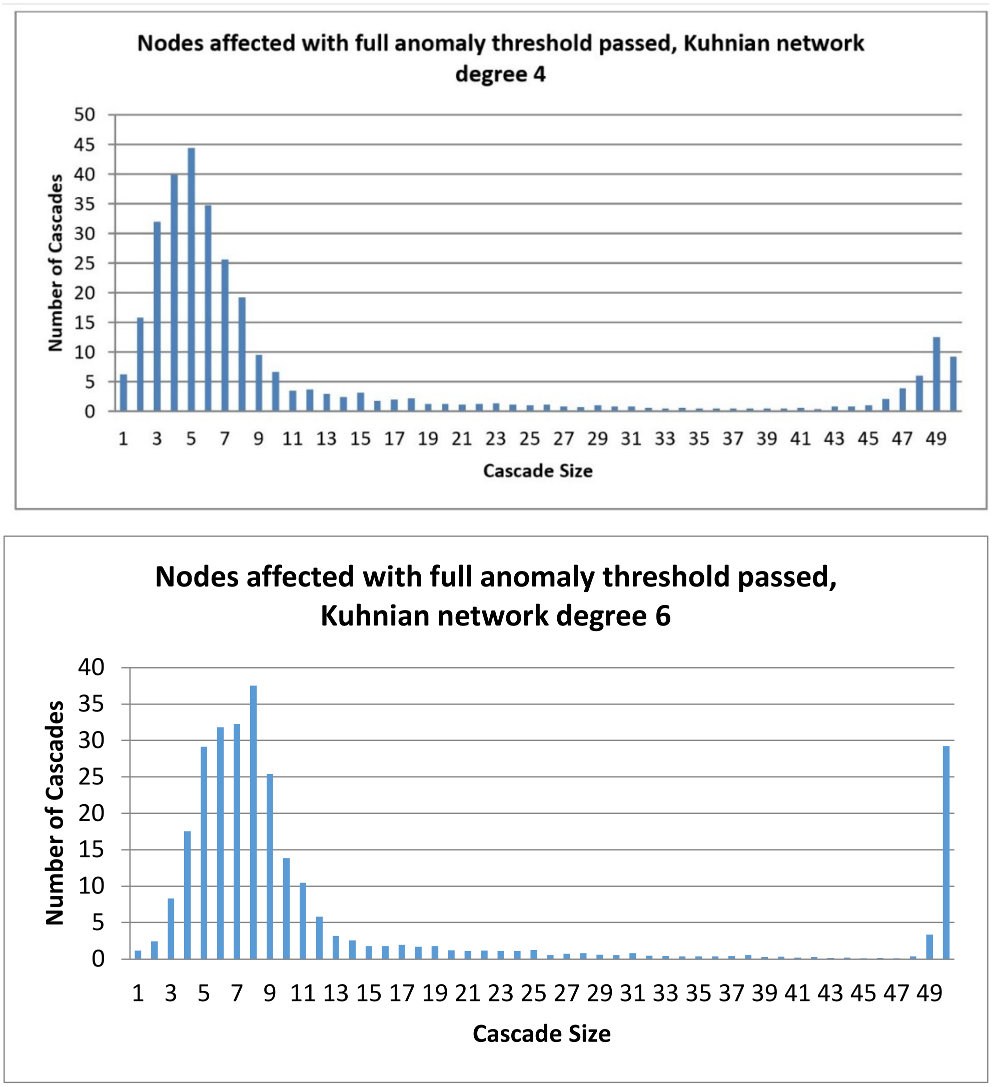

Fig. 8 Cascades of affected nodes in Kuhnian networks, with full anomaly thresholds passed to neighboring nodes.

In this form of the model distributions again start with a power law at low degree, with a cluster that moves slowly right in the pattern of a giant component. At a characteristic degree of 4 and above, however, something importantly different occurs: a significant number of total or nearly total cascades at the extreme upper end. These are cases in which virtually *every* node in the network is affected by a single drop. At a characteristic degree of 6, moreover, those genuinely global cascades of change are among the most common.

IV. Modeling the History of Scientific Change: Popper and Kuhn

We began by noting two aspects of scientific dynamics:. Science grows, and science changes. The first element is an increasing integration of theory over time, modelable as an increasing characteristic degree of a conceptual network as links are added. The second element is a dynamics of scientific change that we have seen to be dependent on degree of integration over time: Popperian falsification cascades or Kuhnian paradigm shifts. By bringing these two aspects together we can create models of science developing, crashing, and rebuilding over time.

Here we take falsification for what it is: the failure of an entire set of links within a sub-network. On the advent of a cascade, we will treat all links involved in the cascade as broken. As the process of scientific development proceeds, however, *new* links will be added.

We start with a Popperian network of 50 nodes with characteristic degree 1, dropping a single falsification somewhere in the network at each iteration. Links from 'falsified' elements are eliminated. But at each iteration an additional random link is added somewhere in the network, representing the force of increased scientific integration. We track the size of integrated scientific theory at a time by simply tracking the number of links in our model of a conceptual structure. The development of science on that Popperian dynamics is shown in Fig. 9.

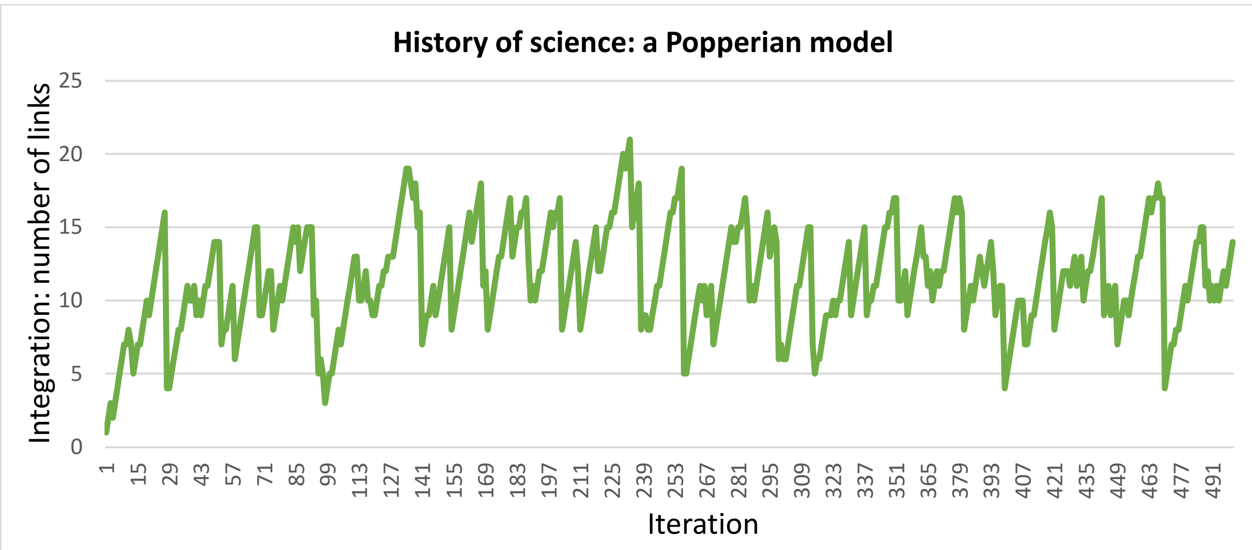

Fig. 9. Scientific change within a Popperian model over time. All links within a falsification cascade are broken, with new links added at each iteration.

Science builds and crashes. On the Popperian picture the crashes are frequent. The integrative links in our model only once peak above 20.

In a similar model for Kuhnian dynamics, starting with same degree and number of nodes, we treat a build-up of anomalies beyond its threshold as 'discrediting' a node. Links from that node disappear. We replace it with a new node, with a new threshold, and again build up links. Figure 10 shows a Kuhnian picture of the development of science.

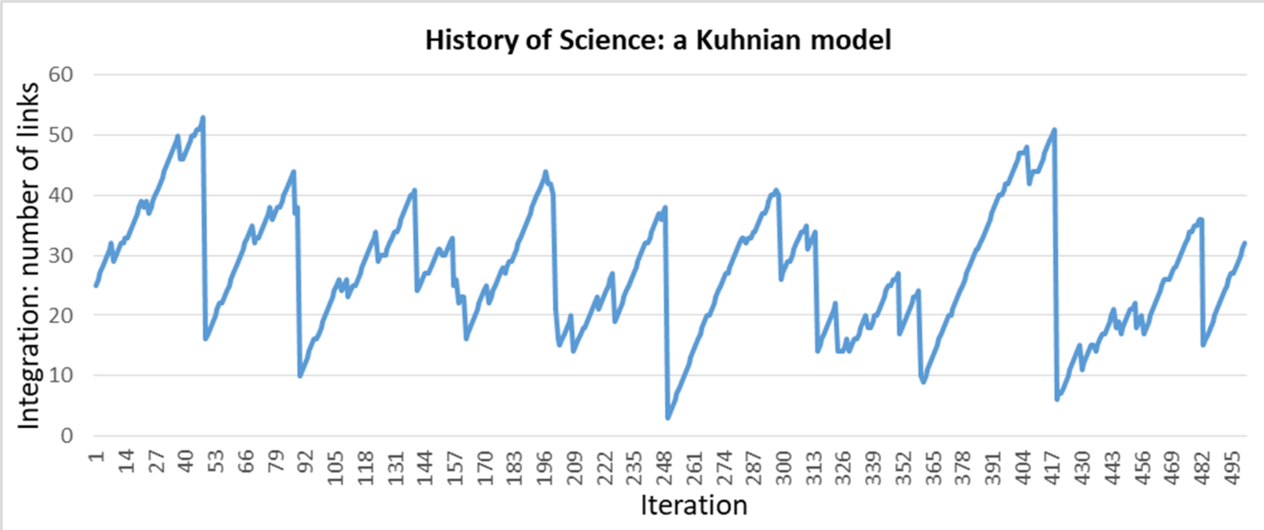

Fig. 10. Scientific change within a Kuhnian model over time, using transfer of a single anomaly. All links from a node that reaches anomaly threshold are broken, with new links added at each iteration.

The first thing to note in these two images is the difference in scale on the y axis. In not to 20 but to 50. The crashes are both less frequent and more drastic.

One might have thought that it would be the Popperian picture of scientific development as conjectures and refutations that would be more rugged. But in fact it is the Kuhnian process that paints a more dramatic picture of progressive accumulations and dramatic crashes—the revolutionary collapses of massive paradigm shifts.

V. Directed Networks: A Second Model of Scientific Paradigms

To this point we have envisaged conceptual systems and scientific paradigms as networks of mutual support between elements. Hence the use of undirected graphs. It can be argued, however, that a more realistic portrayal would use directed graphs instead. We might think of a scientific theory, for example, as a complex of 'if...then' statements between observational phenomena and other conceptual elements. On such a picture our links should be directed: if this holds, then this follows. If this law applies, then this phenomenon is to be expected. Given this as a cause, this can be an expected effect.

Directed network models for conceptual systems have precedent in the long but varied history of conceptual maps. In 1913, John Henry Wigmore developed a 'chart method' for analyzing evidence in a legal case (Wigmore 1913). Wigmore's 'chart method' uses a directed graph (Fig 11). In the 1970s, Robert Axelrod sketched cognitive maps as signed digraphs (Axelrod 1976); generalizations employing continuous values appear in the fuzzy cognitive maps of Bart Kosko (Kosko 1986, 1994, 1997). Cognitive maps as directed graphs have been used extensively as heuristics for group discussion regarding complex issues (Hobbs, Ludsin, Knight, Ryan, Bilberhofer & Ciborowski 2002; van Vliet, Kok & Veldkamp 2010; Soler, Kok, Câmara & Veldkamp 2012; Cakmak, Dudu, Eruygur, Ger, Onurlu & Tonguç 2013; Jeter & Sperry 2013).

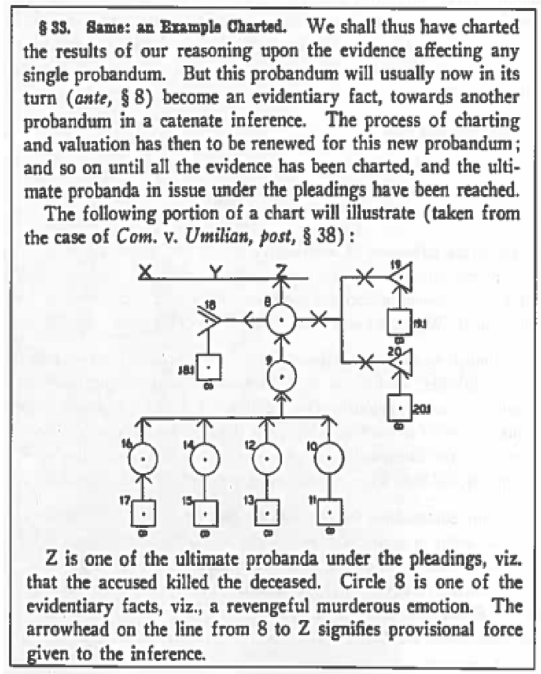

§ 33. Same: an Example Charted. We shall thus have charted the results of our reasoning upon the evidence affecting any single probandum. But this probandum will usually now in its turn (*ante*, § 8) become an evidentiary fact, towards another probandum in a catenate inference. The process of charting and valuation has then to be renewed for this new probandum; and so on until all the evidence has been charted, and the ultimate probanda in issue under the pleadings have been reached.

The following portion of a chart will illustrate (taken from the case of *Com.* v. *Umilian, post*, § 38):

Z is one of the ultimate probanda under the pleadings, viz. that the accused killed the deceased. Circle 8 is one of the evidentiary facts, viz., a revengeful murderous emotion. The arrowhead on the line from 8 to Z signifies provisional force given to the inference.

Fig 11. Wigmore's 1913 outline of legal reasoning in terms of a directed graph

How do the modeling phenomena we've noted play out when our picture of scientific paradigms or conceptual systems in general takes the form of directed as opposed to undirected graphs? Here again the key will be integration of a network measured in terms of characteristic degree. We'll count in-links and out-links together as node degree: average degree is the average of both together.

Figure 12 shows cascade sizes from falsification for directed Popperian networks of increasing degree, with results averaged over 1000 networks.

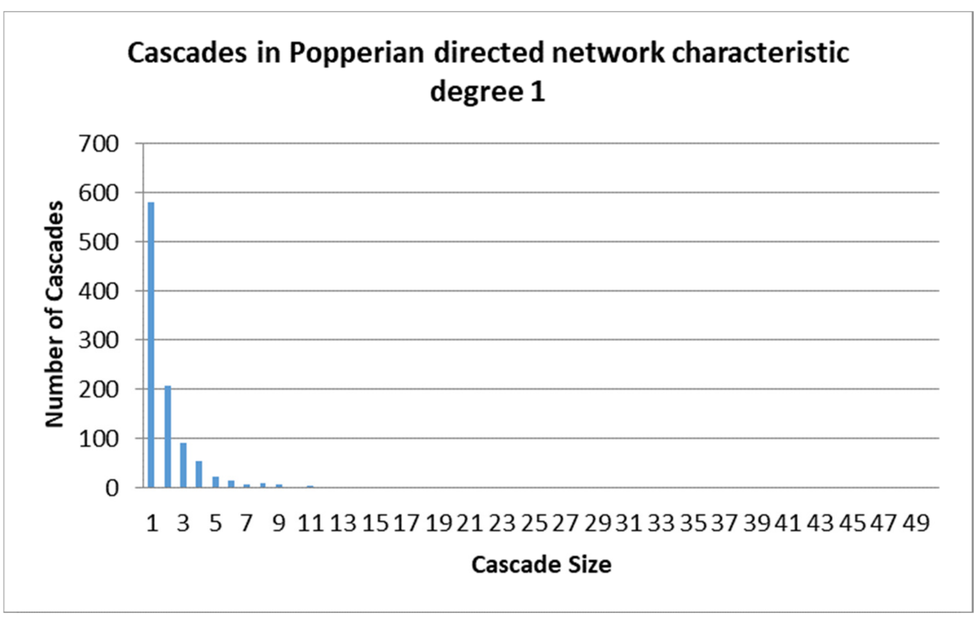

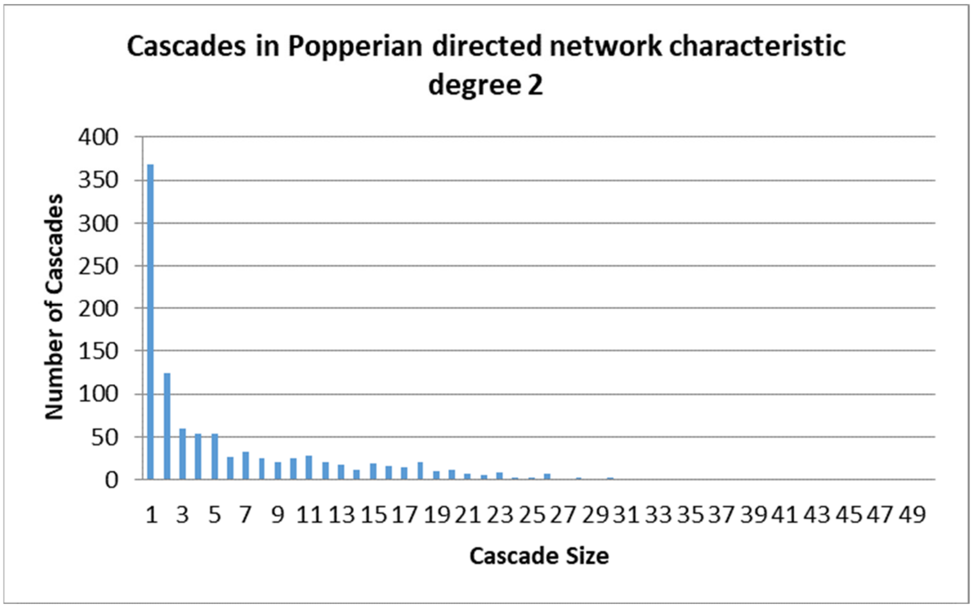
Cascades in Popperian directed network characteristic degree 2
Number of Cascades
0
50
100
150
200
250
300
350
400
1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49
Cascade Size

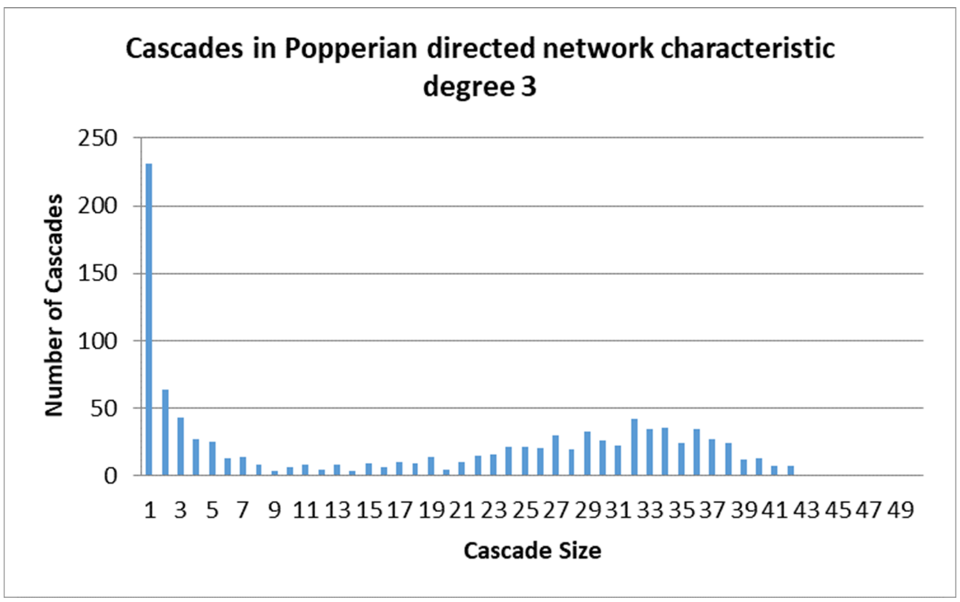
Cascades in Popperian directed network characteristic degree 3
Number of Cascades
0
50
100
150
200
250
1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49
Cascade Size

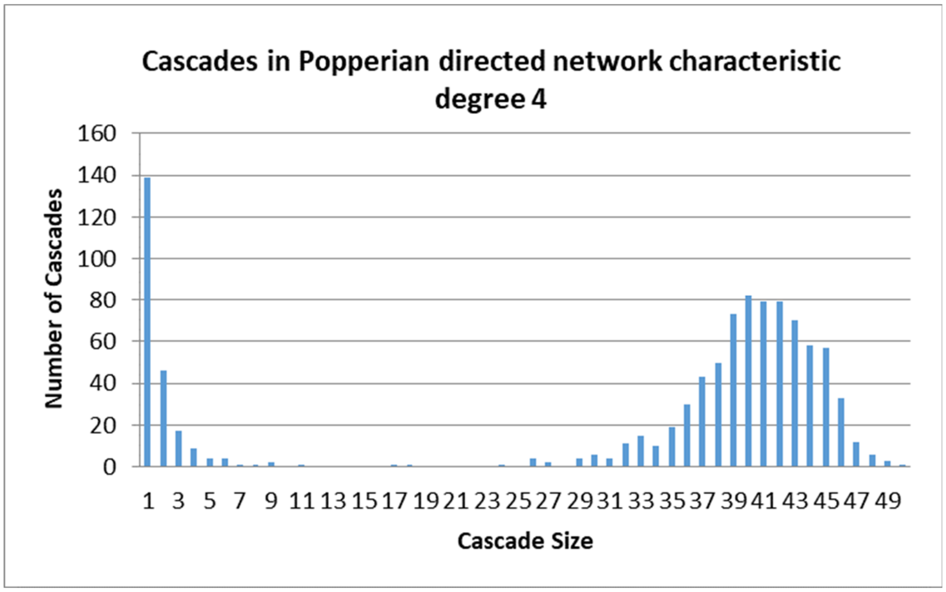

Fig. 12 Distribution of cascade sizes in directed Popperian networks of increasing degree. Note change in scale on y-axis.

The basic Popperian pattern remains the same as in the undirected case: a power-law like distribution at low degree, with the appearance of a clearly bimodal distribution as the integration of a conceptual network increases.

Figure 13 shows the results for a Kuhnian model employing a directed network. We drop 1000 'anomalies' consecutively on nodes with random thresholds between 0 and 5, passing a single anomaly to neighboring nodes when a node reaches its threshold. Cascades are measured as number of nodes fully tipped by a single dropped anomaly.

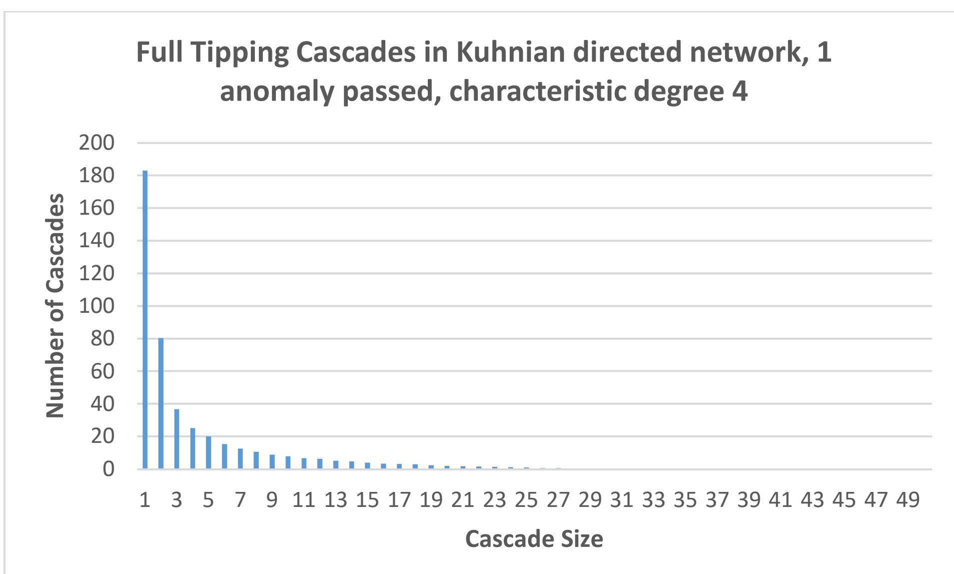

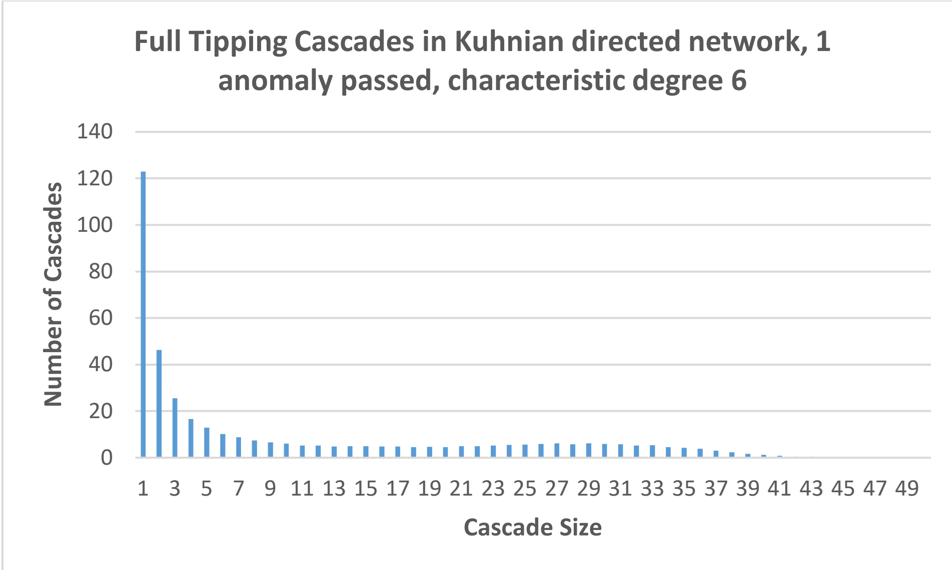

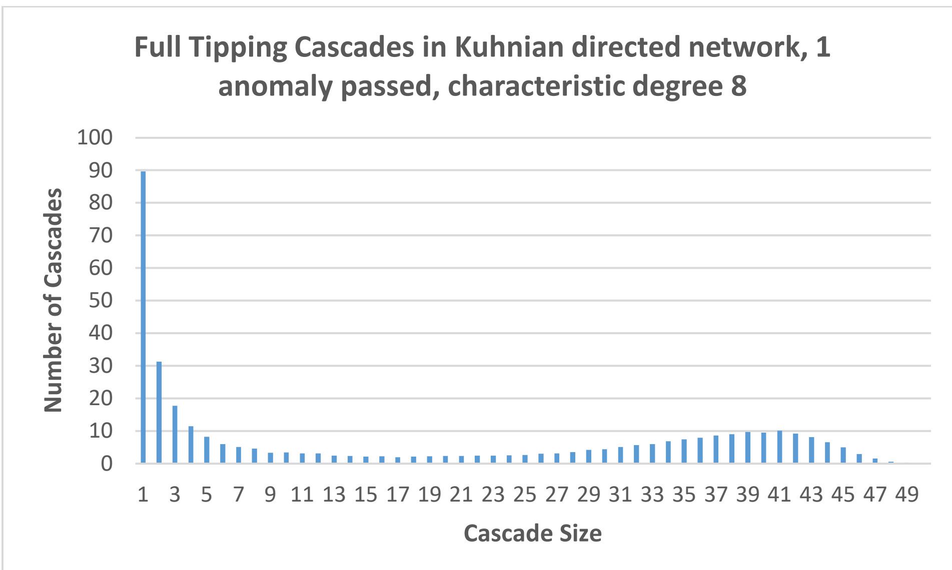

Fig. 13 Distribution of cascade sizes in directed Kuhnian networks of increasing degree. Note change in scale on y-axis.

For both Popper and Kuhn we see the same patterns on both network types: Popperian and Kuhnian dynamics are robust across representations of conceptual networks as either undirected or directed networks.

A formal network analysis in terms of giant components available for the case of directed networks as it was for the case of undirected, though for directed networks one needs to consider not a single giant component but three elements of a ‘bow-tie’ diagram. The core is a ‘strongly connected component’, in which any node has a directed path to every other in the component. In the case of directed graphs, however, we also need to track in-links and out-links.

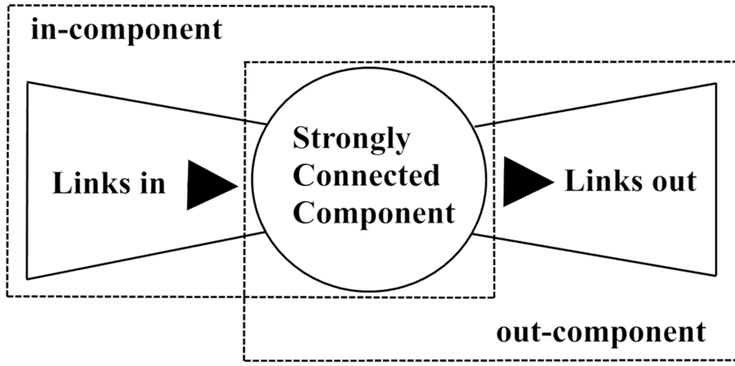

Fig. 14 The bow-tie diagram of components in a directed graph

Despite the fact that we've gone from undirected to directed graphs, the path to a giant component—here a strongly connected component abetted by in- and out- links—is remarkably similar.

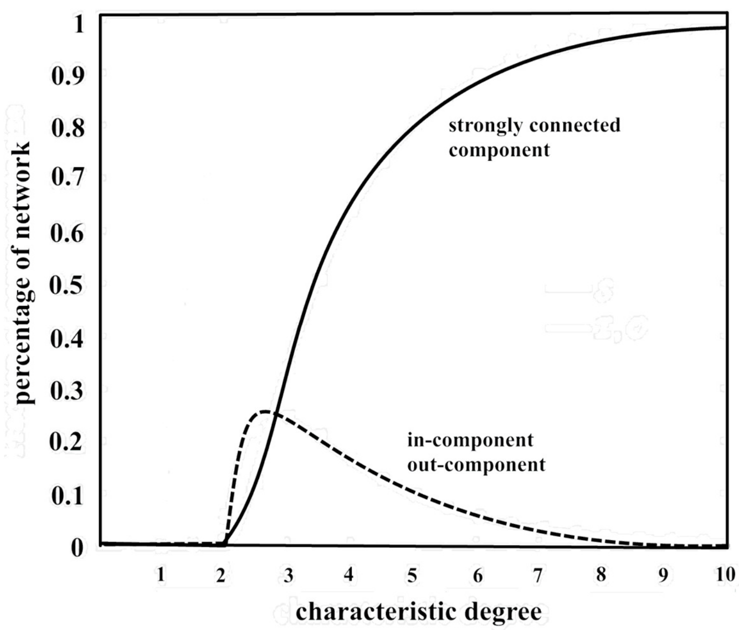

Fig. 15 Formation of strongly connected components, in-component and out-components as percent of a network at different degrees within a directed random network.
Adapted from James B. Glattfelder 2013

The simple message is that here again it is the sudden and dramatic growth of that central component that explains much of the cascade distributions we see in our graphs. Another aspect of this analysis, to which we will return in section VI, is that the proportion of the network occupied by strongly connected components in directed graphs is scale free, independent of network size.

We can also model the developmental history of science within Popperian and Kuhnian models of scientific change (Fig. 16).

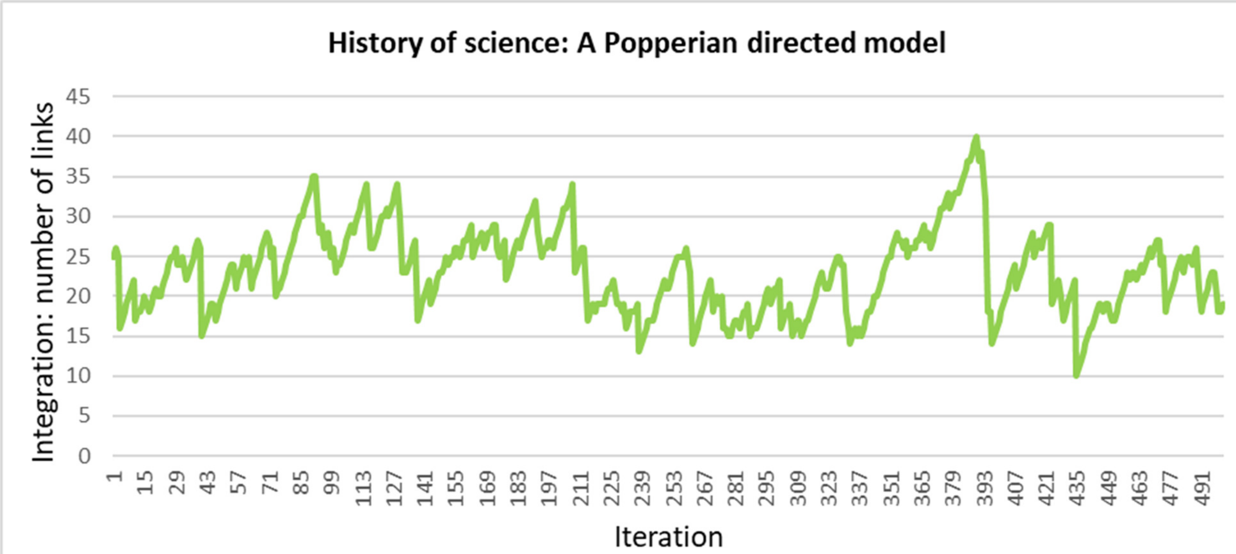

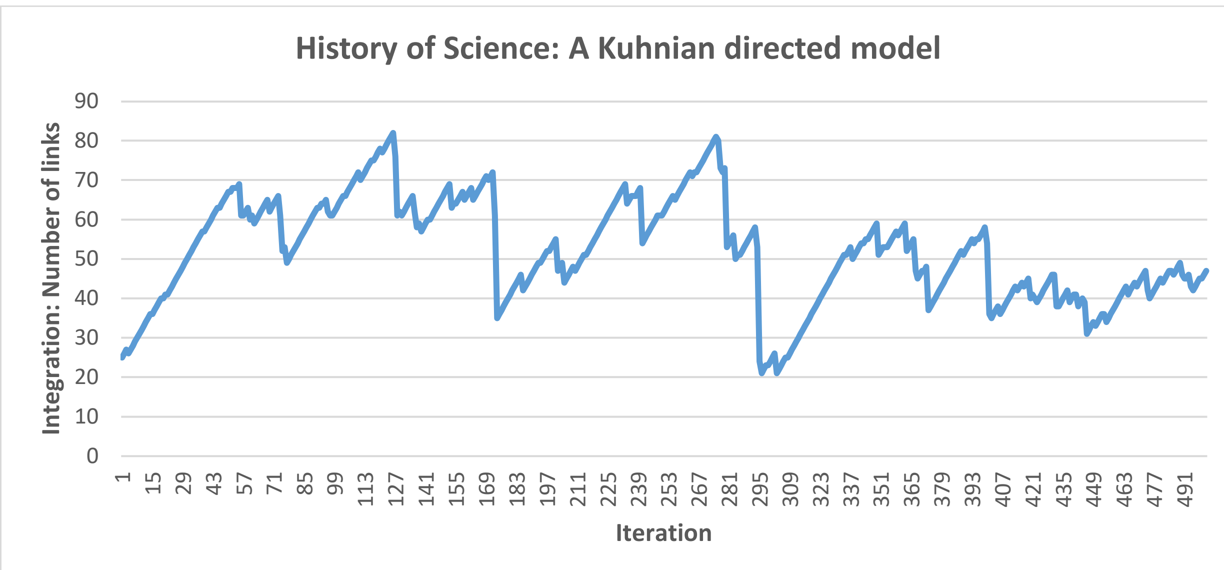

Fig. 16. The history of science in Popperian and Kuhnian models using directed conceptual networks. Beyond the use of directed in place of undirected graphs all parameters are the same as in figures 9 and 10.

One might have thought that we would see a smoother development of models of history in the directed than in the undirected case, since in the latter case cascades must percolate through a more elaborate pattern of directed links. It turns out, however, that the history of science as a series of paradigm shifts is at least as dramatic when conceptual systems are envisaged as directed networks. Here again our results prove robust across this variation in basic modeling of conceptual systems.

VI. Science and Self-organized Criticality

We have one final concept from complex systems to add to the mix: that of self-organized criticality, which finds a new application here.

Self-organized criticality first appears in Bak, Tang & Weisenfeld 1988. In Bak's words, "...many composite systems naturally evolve to a critical state in which a minor event starts a chain reaction that can affect any number of elements in the system." "Large iterative systems perpetually organize themselves to a critical state in which a minor event starts a chain reaction that can lead to a catastrophe" (Bak & Chen 1991).

Avalanches in a sand pile constitute the primary model in Bak's original presentation. Drop a single grain on a sandpile and little or nothing may happen...until something does. Drop a further grain, and a further...and one will have avalanches both small and large. The core idea is that the pile will 'self-organize' toward criticality: without any outside tuning of parameters, the system itself evolves to the point that it's 'ready' for a major avalanche. And it does so again and again.

Self-organized criticality has found wide application both within its original home in physics—with applications in solar, magnetospheric and fusion plasma instabilities—and well beyond. It has emerged as a strong explanatory candidate for patterns of earthquakes, solar flares, and forest fires. It has been proposed as an element of explanation for fluctuations in economic models and for punctuated equilibria in biological evolution (Watkins, Pruessner, Chapman, Crosby & Jensen 2016). Most tantalizing with an eye to conceptual networks and philosophy of science, perhaps, self-organized criticality has both been proposed within the computer sciences as an efficient mechanism of search and within the brain sciences as a crucial mechanism in the functioning brain (Levina, Herrmann & Geisel 2007, 2009; Brochini, Costa, Abadi, Roque, Stofli & Kinouchi 2016; Hoffman & Payton 2018).

What our results seem to hint is that the process of science may be self-organizing as well. Science itself may be an informational instantiation of self-organizing criticality.

It must be admitted that the concept of self-organized criticality has yet no established mathematical formalism or generally accepted definition. Even in Bak's original presentations, the concept is outlined not by strict definition but in terms of 'marks' of self-organized criticality

One of those marks is 'flicker noise' or 1/f noise. White or random noise shows no correlation from point to point. In flicker noise, ubiquitous in natural systems, there *is* a strong correlation between points and their predecessors—a clear indication of a path-dependent dynamics. Correlation of that type is clear throughout the phenomena we've tracked in both Popperian and Kuhnian networks.

Bak's other major mark of self-organized criticality is the fractal characteristic of scale-invariance. In theory, the relative size of cascades should be the same in sand piles regardless of size. That characteristic is also clear here.

Our model results throughout have used networks of 50 nodes. But they didn't have to. Beyond a critical point, our results scale up regardless of the size of the network. Figure 17 compares cascades on Popperian networks of 50 and 100 nodes. Although more finely tuned in the second case, with results across more options for cascade sizes, the same clear pattern appears.

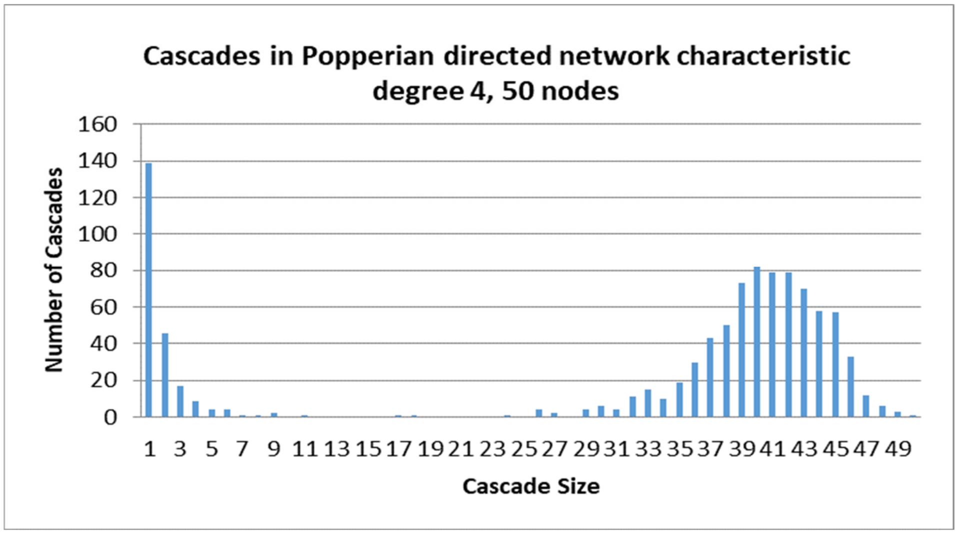

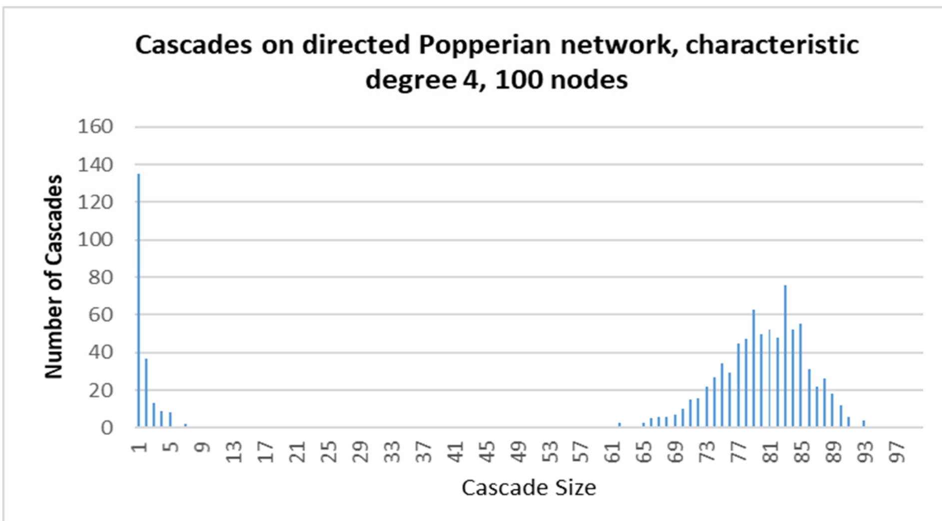

Fig. 17 Comparative cascade sizes in Popperian networks of 50 and 100 nodes

As we've indicated, a core explanation for the cascade patterns tracked here is in the emergence of giant components in networks, both undirected and directed. It is because giant components of a particular size appear—tied directly to the characteristic degree of a network—that cascades of that size dominate our graphs. What is of importance here is the proportion of a network occupied by a giant or strongly connected component, in the case of either undirected or directed graphs, is scale-invariant. Beyond a surprisingly low phase transition, the proportion of a network occupied by such a component is independent of the size of the network. To the extent that our cascade distributions are tied to the presence of such components, precisely because the network proportion of those components is scale invariant, our cascade distributions will be as well.

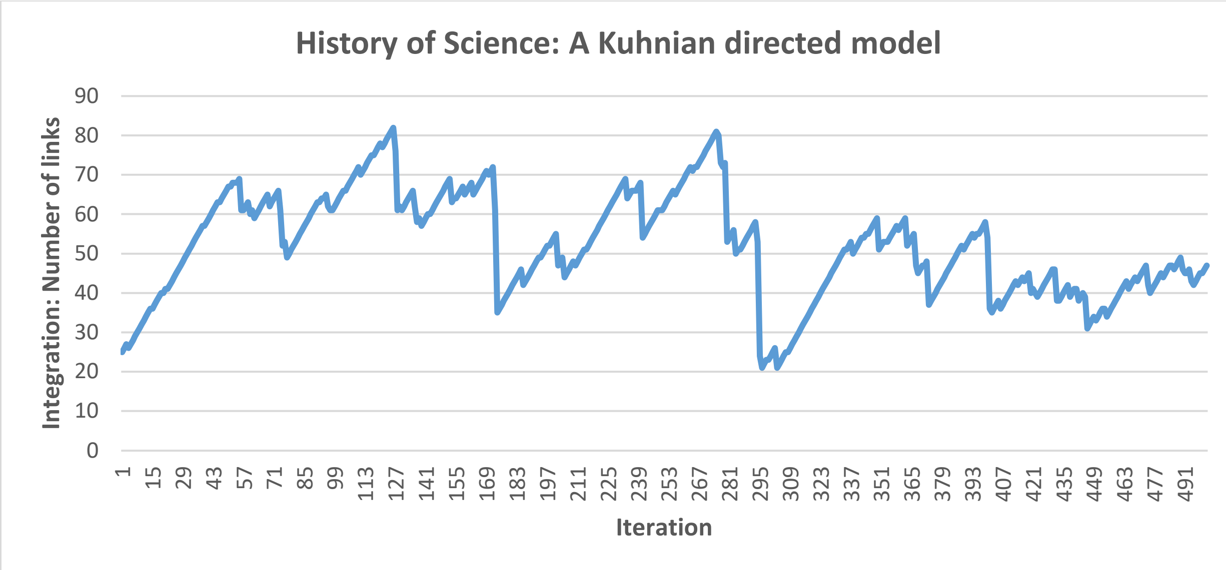

Fig. 18 The history of science on a Kuhnian directed model.

A clearer example of self-organization than the Kuhnian dynamics of scientific change that we've modeled would indeed be hard to find. The system self-organizes towards that point at which a major cascade across much of the network—a revolutionary paradigm shift—is well-nigh inevitable.

VII. Conclusion: Prospects for Expanding the Models

It must be admitted that this remains a work in progress. There are three clear directions for future development.

The models for conceptual networks used here have been the simplest: undirected and directed random graphs. One clear direction for future work is to investigate cascade phenomena using other network types as representations of conceptual systems. One prime candidate is Boolean networks. Do dynamics change when we make explicit a structure developed in terms of logical connectives? Another clear candidate is Bayesian networks. What are the characteristics under which conceptual cascades occur in networks within which nodes update priors on input from other nodes? Those remain unanswered questions.

There is another natural expansion of the model that is clearly called for. Here we have followed Kuhn and Popper in treating a conceptual system as a shared paradigm, a single possession of a scientific community. But of course different individuals have different conceptual systems. Given specific patterns of communication, certain changes in one may effect certain changes in others. Cascades can be expected to happen on the social level as well. What is called for in expansion is a two-level model embedding individual conceptual networks within a second level of social communication.

A third direction called for—here and in agent-based modeling in general—is a closer link to empirical data. One aspect of our results has been a picture of the history of science on Popperian and Kuhnian models. The work offered here remains very much on the abstract end, the theoretical rather than the empirical, philosophy of science rather than history of science. The question of whether the history of science has something like the topography of our graphs remains an open question, empirical and hard.

References

Bak, P., C. Tang & K. Weisenfeld (1988) "Self-Organized Criticality," *Physical Review A* 38 (1): 364-374.

Bak, P. & K. Chen (1991). Self-organized Criticality." *Scientific American*, 264 (1): 46-53.

Bovens, L. & S. Hartmann (2003). *Bayesian Epistemology*. Oxford: Oxford Univ. Press.

Brochini, L., A. de A. Costa, M. Abadi, A. C. Roque, J. Stofli & O. Kinouchi (2018). "Phase gransitions and self-organized criticality in networks of stochastic spiking neurons," *Scientific Reports* 6: e5831. Doi: 10.1038/srep35831

Cakmak, E. H., H. Dudu, O. Eruygur, M. Ger, S. Onurlu & Ö. Tonguç (2013). "Participatory fuzzy cognitive mapping analysis to evaluate the future water in the Seyhan Basin." *Journal of Water and Climate Change* 4 (2) 131-145.

Carnap, R. (1969). *The Logical Structure of the World*. Berkeley: Univ. of California Press.

Carnap. R. (1966). *Philosophical foundations of Physics: An Introduction to Philosophy of Science*. New York: Basic Books.

Curd, M. V. (1982). "The Rationality of the Copernican Revolution." *PSA: Proceedings of the Biennial meeting of the Philosophy of Science Associaiton*, Vol. 1982, 3-13.

Erdös, P. and A. Rényi (1959). On random graphs. *Publicationes Mathematicae* 6, 290-297.

Glattfelder, J. B. (2013). *Decoding Complexity: Uncovering Patterns in Economic Networks*. New York: Springer.

Goodman, N. (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard Univ. Press.

Hempel. C. (1966). *Philosophy of Natural Science*. Upper Saddle River, NJ: Prentice-Hall

Hobbs, B.F., S. A. Ludsin, R. L. Knight, P. A. Ryan, J. Bilberhofer & J. J. H. Ciborowski (2002). "Fuzzy cognitive mapping as a tool to define management objectives for complex ecosystems." *Ecological Applications* 12 (5): 1548-1565.

Hoffman, H. & Payton, D. W. (2018) "Optimization by Self-Organized Criticality." Scientific Reports 8, article number doi:10.1038/s41598-018-20275-7

Jeter, A. J. & R. C. Sperry (2013). "Fuzzy Cognitive Maps for Product Planning: Using Stakeholder Knowledge to Achieve Corporate Responsibility." 46th Hawaii International Conference on System Science, 925-934.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Kuhn, T.S. (1969). "Postscript-1969." In *The Structure of Scientific Revolutions*, Fourth Edition. Chicago: University of Chicago Press.

Kuhn, T. S. (1970). "Logic of Discovery or Psychology of Research?" In I. Lakatos & A. Musgrave, eds., *Criticism and the Growth of Knowledge*. Cambridge: Cambridge Univ. Press.

Lakatos, I. (1968). "Criticism and the Methodology of Scientific Research Programmes", *Proceedings of the Aristotelian Society*, 69: 149–186.

Laudan, L. (1978). *Progress and its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.

Levina, A., J. M Herrmann, & T. Geisel (2007). "Dynamical synapses Causing Self-Organized Criticality in Neural Networks," *Nature Physics* 3, 857-860.

Levina, A. J. M Hermann & T. Geisel (2009). "Phase transitions towards criticality in a neural system with adaptive interactions," Phys. Rev. Lett 102 (11): 118110.

Newman, M. E. J. 2010. *Networks: An Introduction*. Oxford: Oxford University Press.

Perrin, J. (2005). *Brownian Motion and Molecular Reality*. New York: Dover.

Popper, K. (1959). *The Logic of Scientific Discovery*. Oxford: Basic Books.

Popper, K. (1963). *Conjectures and Refutations*. New York: Routledge & Kegan Paul.

Quine, W. V. O. (1951). "Two Dogmas of Empiricism." Philosophical Review 60 (1): 20-43.

Quine, W. V. O. (1953). *From a Logical Point of View*. Cambridge MA: Harvard University Press.

Quine, W. V. O & J. S. Ullian (1978), *The Web of Belief*. New York: Random House.

Soler, L.S., K. Kok, G. Câmara & A. Veldkamp (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.

van Vliet, M., K. Kok and T. Veldkamp (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.

Watkins, N. W., G. Pruessner, S. C. Chapman, N. B. Crosby & H. J. Jensen, "25 years of self-organized criticality: Concepts and Controversies," *Space Science Review* (2016) 198: 3-44.