Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence

Patrick Grim, Daniel J. Singer,* Steven Fisher, Aaron Bramson, William J. Berger, Christopher Reade, Carissa Flocken, and Adam Sales[1]

Abstract

A scientific community can be modeled as a collection of epistemic agents attempting to answer questions, in part by communicating about their hypotheses and results. We can treat the pathways of scientific communication as a network. When we do, it becomes clear that the interaction between the structure of the network and the nature of the question under investigation affects epistemic desiderata, including accuracy and speed to community consensus. Here we build on previous work, both our own and others', in order to get a firmer grasp on precisely which features of scientific communities interact with which features of scientific questions in order to influence epistemic outcomes.

Here we introduce a measure on the landscape meant to capture some aspects of the difficulty of answering an empirical question. We then investigate both how different communication networks affect whether the community finds the best answer and the time it takes for the community to reach consensus on an answer. We measure these two epistemic desiderata on a continuum of networks sampled from the Watts-Strogatz spectrum. It turns out that finding the best answer and reaching consensus exhibit radically different patterns. The time it takes for a community to reach a consensus in these models roughly tracks mean path length in the network. Whether a scientific community finds the best answer, on the other hand, tracks neither mean path length nor clustering coefficient.

---

* Patrick Grim and Daniel J. Singer contributed equally to this research. For inquiries, please contact Patrick Grim (patrick.grim@stonybrook.edu) or Daniel J. Singer (singerd@phil.upenn.edu).

[1] Patrick Grim is Distinguished Teaching Professor in Philosophy at Stony Brook and Visiting Professor with the Center for Study of Complex Systems at the University of Michigan. Grim's current work centers on computational modeling and philosophical logic.

Daniel J. Singer is an Assistant Professor of Philosophy at the University of Pennsylvania. He primarily researches the nature of epistemic normativity and uses computational models to explore the social structure of inquiry.

Steven Fisher is a graduate of the University of Michigan in Complex Systems. He is currently working with early stage startups to optimize technology strategies.

Aaron Bramson is a visiting scientist in the economics department of Gent University and at the Riken Brain Science Institute. His research focuses on complex systems methodology: developing new tools for building and analyzing agent-based, network, and mathematical models.

William J. Berger is a doctoral candidate in the department of political science and NSF IGERT Fellow at the University of Michigan where he studies the institutional aspects of affective trust.

Christopher Reade is a consultant at ThoughtWorks Inc. and a graduate of the Gerald R. Ford School of Public Policy at the University of Michigan.

Carissa Flocken studies Complex Systems at the University of Michigan. As a research assistant at the Santa Fe Institute, she explores the evolution of social network topology using computational methodology.

Adam Sales is a postdoctoral researcher at Carnegie Mellon University. He studies statistical methods for evaluating educational programs using observational data and likes hanging around philosophers.

I. Introduction

No scientist is an island. Scientists are connected to other scientists in a myriad of ways. Data and results are shared via journals. New ideas, hypotheses, and hunches are communicated formally at conferences and informally at casual gatherings. Scientists share graduate students, lab space, statistical techniques, and research-related rumor; they are networked by personal contact, phone, letter, text and e-mail. We investigate the role that these communication pathways play on the quality of scientific results.

From previous work we know that a scientific community may learn more when its individual scientists share less (Grim 2006, 2007, 2009a, 2009b; Zollman 2007, 2010a, 2010b, forthcoming). In terms of the scientific goal of producing true theories alone, increased communication among scientists may not always be a good thing. Seventeenth century science was characterized by distributed informational networks with limited linkages between investigators. Twenty-first century science is characterized by much more highly connected networks, in part due to the internet. For some scientific goals, targets, and aspects of investigation, the network structure of 17th century science appears to be superior to our own.

Here, we focus on the more specific question of how communication networks interact with the kinds of problems that scientists confront in ways that either further or inhibit scientific progress. A complete answer to this question would require attention to the complications of funding structures, human psychology, laboratory logistics, and professional and personal allegiances, among many other aspects of real science. We aim to make some progress on this question without those complications by considering abstract and idealized scientific communities. We treat scientists as primarily concerned with truth-seeking and make many idealizing assumptions about science and the psychology of scientists. By means of abstract modeling, we hope to get at least an abstract grasp on some major components of the messier problem.

We use agent-based computer simulations to explore the relationship between communication networks and aspects of scientific projects, while making the modeling assumptions as explicit as possible. We represent the scientific pursuit as the attempt to find the highest peak on an epistemic landscape. We then construct a measure of problem difficulty for those landscapes, which we use to analyze how the interaction between the investigators' communication network and the problem they're investigating influences scientific success, i.e. whether any scientist finds the peak. We also use the models to investigate the formation of scientific consensus and examine the trade-offs between the goals of success and consensus in terms of the difficulty of the particular question at issue.

II. Epistemic Landscapes, Problem Difficulty, and the Fiendishness Index

Some research questions are easier to answer than others. However, what counts as a difficult problem varies from field to field, and attempts to codify and measure problem difficulty vary accordingly. In computational complexity research, a well-known line is drawn between P problems, whose answers are computable in polynomial time, and NP problems, answers for which are only verifiable in polynomial time. That is one measure of question difficulty. It also remains a notably hard open question whether P = NP. In optimization research, there are also attempts at finding a satisfactory measure for fitness functions resistant to solution by genetic algorithm and evolutionary computation. Isolation, fitness-distance correlation, fitness variance, and epistasis variance have all been proposed as measures of difficulty in this regard, but for each there are some intuitively easy cases that those measures improperly classify as *hard* and some intuitively hard cases they class as *easy* (Naudts & Kallel 2000; He, Reeves, Witt & Yao 2007). The term 'wicked problem' has been applied in social policy with a set of fairly loose criteria. Novel or unique problems are said to be wicked if they have no definitive formulation, definitive solution, or even a definitive set of alternative solutions. A problem is wicked if it is a one-shot problem for which an answer cannot be generalized from other cases or generated by trial or error, and the investigation of which carries no "stopping rule" (Rittel & Webber 1973; Conklin 2003).

The modeling framework we present below is flexible enough to model many types of problem difficulty. In this first paper we restrict ourselves to problems that are difficult for the straightforward empirical reason that their solution may take a lot of trial-and-error to find; e.g., practical medical effectiveness, true causal pathways, how a specific protein folds, or where a treasure is buried. Some questions of this kind prove difficult because of tempting answers that are merely second- or third-best. Other questions prove particularly difficult because the answer (e.g., the best treatment, the true causal chain, the correct folding pathway, or the actual treasure location) is a "needle in a haystack," hard to find because it is hidden in a tiny portion of the hypothesis space. For this kind of question, the shape of the epistemic terrain makes it difficult to find the right answer. We'll call a difficult problem of this type 'fiendish.'

In our models, graphs are used to represent particular aspects of the epistemic structure of a problem. For a question about which dosage of an antiviral drug is most effective, hypothesized best dosages are represented on one axis. The other axis gives a measure of how good each hypothesis is—the percentage of patients who show remission at a given dosage, for example. These *epistemic landscapes* are similar to fitness landscapes in biology except that the height indicates epistemic fitness of a hypothesis rather than ecological fitness of a genotype.[2] We think of the epistemic landscape as giving an objective picture of the problem under investigation: it maps every possible hypothesis to the strength of the epistemic upshot of the deliverances of the world for a research investigating that hypothesis, which we can think of as the epistemic goodness of that hypothesis.

The epistemic landscape gives a 'God's eye view' of the epistemic payoffs of the entire collection of hypotheses. All the individual researcher has access to is a small collection of particular points of investigation on that landscape. A researcher adopts a hypothesis by taking a value on the x-axis. Only then can he find the epistemic payoff for that hypothesis represented in the landscape. Researchers are doing better when the epistemic payoff of their hypothesis is higher, so researchers aim to adopt the hypothesis with the highest payoff. What makes a question fiendish is that the epistemically best hypothesis is hard to find. Progressively more fiendish questions are represented by the epistemic landscapes in Fig. 1. In our models, we restrict ourselves to linear representations to reduce the complexity of the model; in many cases, such as combination of two cancer treatments, an epistemic landscape will be better represented in more dimensions (Fig. 2) Empirical questions may be difficult in other regards as well (e.g. due to ambiguous, missing, or conflicting evidence), but our inquiry here is focused on fiendishness.

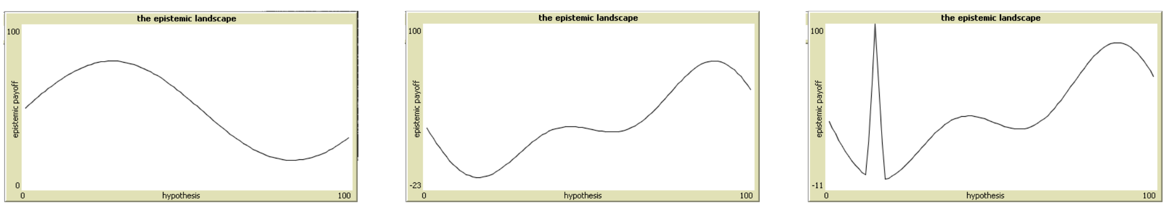

(a) (b) (c)

Fig. 1 An abstract representation of progressively more fiendish empirical questions

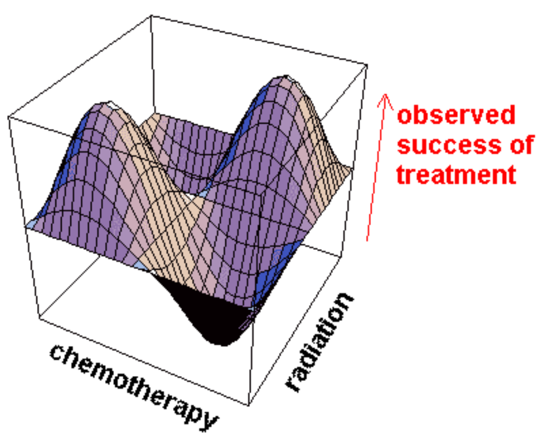

Fig. 2. An idealized epistemic landscape for effectiveness of combined cancer therapies

As the examples make clear, landscapes can contain both a global maximum peak and one or more sub-optimal local maxima. The number, location, height, and position of the maxima of the landscape can affect its

[2] Weisberg and Muldoon (2009) use epistemic landscapes in which 'fitness' is taken as scientific significance.

fiendishness. Intuitively, finding the needle in the haystack is a more fiendish problem than finding a medium-sized object in the haystack because a searcher is less likely to happen upon the needle in the course of the search. We provide a formal measure of fiendishness that is meant to capture this idea: The *fiendishness index* of a landscape, relative to a number *N*, is the probability that for *N* random points on the landscape, the highest point will not be within the basin of attraction of the global maximum. The basin of attraction for a maxima is the set of points on the landscape for which simple hill-climbing will result in that maxima. Lower probabilities that the highest random point is within the highest peak's basin of attraction correspond to higher fiendishness. Alternatively stated, higher fiendishness scores correspond to landscapes on which there is a lower probability that hill-climbing from the highest randomly chosen point results in the landscape's global maximum.

There will be uncountably many different landscapes that have any given fiendishness index. In what follows, we focus on simple landscapes constructed in terms of just two peaks: a global maximum that takes the form of a "needle in a haystack" spike, with a height (success level) of .95, and a local maximum hump with a height of .85 (Fig. 3).

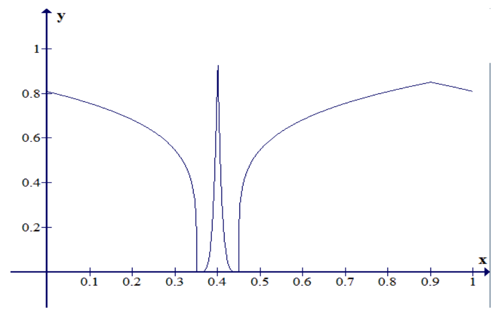

Fig. 3. A simple "spike and hump" epistemic landscape.

We construct our landscapes using a formula in which we can vary the curvature of the spike and hump, without varying the width of their bases. For the spike (shown between x = 0.35 and x = 0.45 in Fig. 3),

$$f(x) = 0.95\left(1 - 2\left(abs\left(10x - 0.5\right)\right)\right)^{z}.$$

For the hump,

$$f(x) = 0.85\left(1 - 2\left(abs\left(\frac{10}{9(x - 0.1)} - 0.5\right)\right)\right)^{\left(\frac{1}{z}\right)}.$$

The variable z is used in both to adjust the curvature of the spike and hump. The fiendishness index tracks the probability that the highest of *N* random points on the landscape will be in the basin of our global maximum. Using the formulae above, we can also think of fiendishness in terms of the 'fiendishness exponent' z. Figure 4 shows the increasingly fiendish landscapes generated with these formulae for exponents 1, 2, 3, 5, and 8. Here for simplicity we exhibit the spike to the left (which does not affect the fiendishness index).

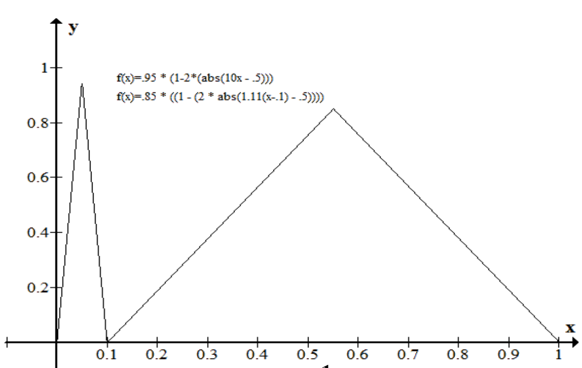

z = 1

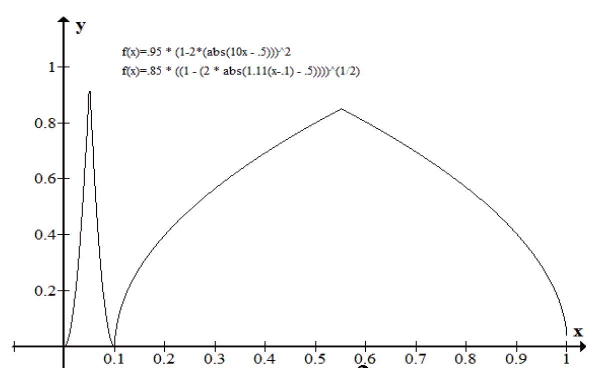

z = 2

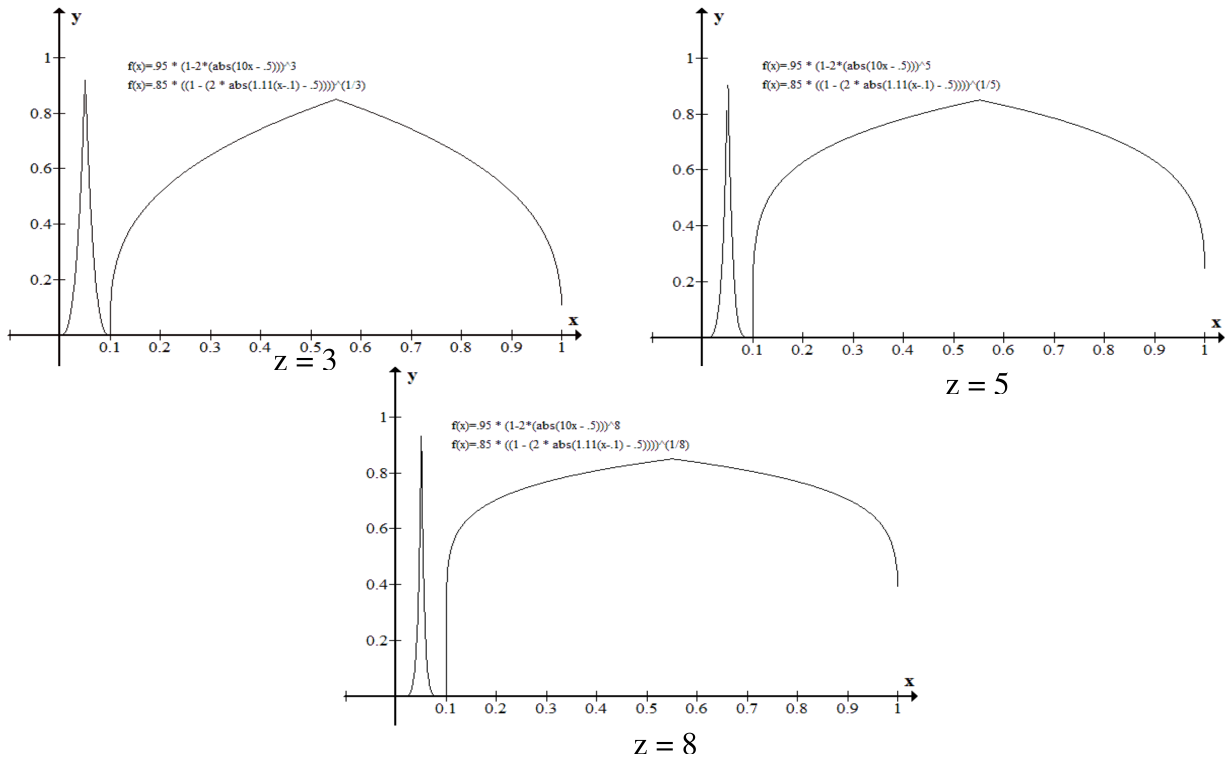

Fig. 4. Epistemic landscapes for exponents z = 1, 2, 3, 5, and 8 showing increasing fiendishness.

Though the fiendishness exponent determines the topology of a particular class of landscapes, the particular probability of the highest of *N* random points falling in the basin of the global maximum will be relative to *N*. In other words, how improbable it will be for the highest random drop to land in the basin of the global maximum will depend on the number of random drops we are allowed. Figure 5 shows the probability of the highest drop landing in the spike area for *N* between 1 and 100 and for fiendishness exponents of z = 1, 3, 5 and 7. In what follows, we will focus on the comparative success of different scientific networks with the same number of agents (fixed value of *N*) across a range of landscapes. We vary the fiendishness index of the landscapes through changes in the parameter z, which tracks the fiendishness index when the number of agents is held fixed.

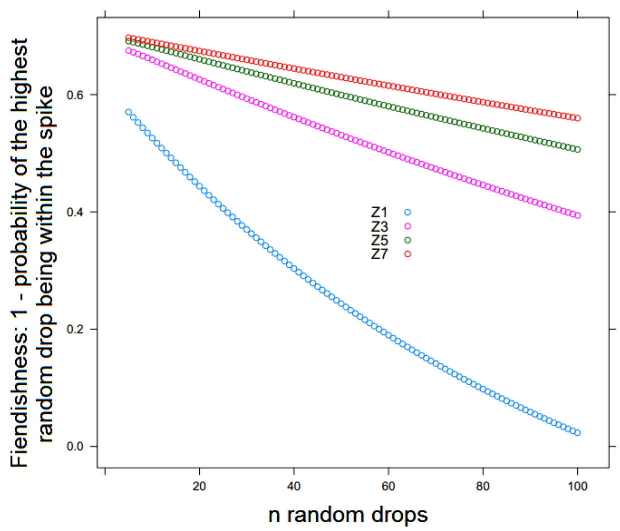

Fig. 5. (1 - fiendishness index) as a function of the parameter z and the number of random drops.

III. Modeling the Social Structure of Science

Using our epistemic landscapes, we can produce a highly idealized, simple simulation model of the social structure of a population of scientists. We consider collections of scientist agents, each of whom has a hypothesis, which we model as the agent having a location on the x-axis of the landscape. Agents are also connected to other agents. At each round of the model, agents receive feedback from the landscape and other agents, and then they change their hypothesis based on the feedback. The feedback consists of the epistemic fitness (the y-value) of their current hypothesis and the hypotheses of those with whom they are connected.

For simplicity, the connections between agents in the models presented here are bi-directional and do not change over time. The basic modeling framework could easily be extended to allow for models with dynamic and directed networks, networks in which different information is shared, networks in which feedback is received without perfect fidelity, and models with multiple networks used at different times and for different purposes (e.g., one for colleagues, another for conferences, and another for scientists who read published research papers).

Our models differ from some related models in important ways. The models of Weisberg and Muldoon (2009) and Hegselmann (2013) put agents on a checkerboard landscape and allow them to see only the epistemic fitness of the hypotheses bordering their current location. By doing so, the models limit the communication of agents to those near them in the hypothesis space. We believe that to model the social nature of science, models must allow the social network structure (the communication network, in this case) to vary independently of the epistemic landscape. Scientists often do consult with others who either moderately or severely disagree with themselves, and the information gained in such a consultation may influence that agent's later views. Our modeling framework allows us to model communication between agents independently of the similarity of their hypotheses.

Our models also differ from the models of the social structure of science that use bandit problems (e.g. Zollman 2007, 2010a). In those models, agents can be thought of as playing mixed strategies on a collection of slot machines. Agents receive binary (win or lose) data about the machine they play on each round and update their hypotheses about which machine is the most advantageous. These models do allow the network structure to come apart from the epistemic landscape, but the epistemic landscape itself is too simple. There are a relatively small number of hypotheses that agents explore (one for each slot machine), and the hypotheses at issue do not bear any interesting topological relationship to each other, like the hypotheses in our epistemic landscapes do. Along with Weisberg and Muldoon (2009), we think that there is a natural and important aspect of hypotheses captured by the topological relationship imposed on them by the landscape, for example that of some hypotheses being "nearer to" or "further from" others. When considering problems like the effectiveness of combined cancer therapies, we can model agents as moving to nearby hypotheses in some circumstances (e.g. when the agent receives some evidence that less radiation is more effective), but we can also model agents radically changing their opinions by moving to more remote hypotheses when that is appropriate. Bandit problem models don't offer an analogous notion of closeness. A benefit of that approach is that it can model agents as being unsure about the successfulness of their current hypotheses, but our models can easily be extended to model that uncertainty as well.

In general, models of the social structure of science like ours consist of three important parts. First, an *epistemic landscape* provides, for individual agents (or individual scientific labs), a connection between the agent's epistemic position—her hypothesis or project—and the epistemic outcome associated with that position. Second, these models contain a *social element*, which in our case consists of the agents' communication network. Similar models might also include other social elements, such as funding sources, allegiances due to training, and journal distributions. The final element of these models is bridge between the first two elements: an *updating rule*, which informs agents about what their future epistemic position should be in light of their current epistemic position and the information provided by the social element of the model.

A potential objection to the models of Weisberg and Muldoon (2009) and Hegselmann (2013) is that they don't distinguish between the first two elements. A potential objection to the bandit problem models of Zollman (2007, 2010a) is that they have too meager an epistemic landscape to be compatible with natural ideas about the updating rule that bridges the social aspects of the model and the epistemic landscape.

IV. Models of Epistemic Success

In the models we present here, we start with a population of agents, each of whom starts with a randomly-assigned hypothesis, represented by a single value between 0 and 1 on the x-axis of an epistemic landscape. In testing their hypotheses, our agents accumulate data on their success as feedback — a percentage of patients cured, for example. Our agents can see not only the success rate of their own hypothesis but also the success rate for the hypotheses of those to whom they are linked. Agents then change their hypotheses based on their own success rate and the success rates of those to whom they are linked. If a contact's hypothesis is better supported by the data than yours or any of your other contacts' hypotheses, you shift your hypothesis toward that hypothesis.

Crucial to this model is an element of speed, which may be seen as either a measure of conservatism or of the limitations of neighbor influence. In shifting toward what appears to be a more successful hypothesis, does an investigator jump immediately to that hypothesis, or do training, prior commitment, familiarity, remaining doubts and inertia dictate a move that is only some portion of the distance in that direction? We build in a simple form of the latter assumption by having agents move merely halfway toward a superior hypothesis at any step. More complex models might treat speed as a function of the hypotheses in play, their epistemic success, or even the past success of the agent.

The model also builds in a "trembling hand": in moving toward another's hypothesis, you may be slightly off, since your lab conditions may be slightly different from that of the other agent, you might slightly misunderstand the hypothesis, or your samples may be slightly biased. In earlier models we used a margin of error of landing randomly within an x-distance of 0.04 (Grim 2006, 2007, 2009a, 2009b). Here we use a trembling hand that is a function of the distance from a point to its target hypothesis. An agent approaches a perceived target hypothesis within a trembling hand window, centered on the target, that is 10% of the distance to that target. Farther targets have a wider window of approximation. Closer targets have a narrower window. At a closer distance, in other words, an agent is better able to "see" the target hypothesis and thus better able to approximate it. Using an intermediate speed and a margin of error results in a greater exploration of the epistemic landscape (Lazer & Friedman 2007).

We consider several topologies for the hypothesis axis of the epistemic landscape below, including one that has no endpoints in the hypothesis space at all. For those cases in which a target hypothesis is close to the edge of a landscape, with the possibility that the agent might overshoot the hypothesis space by adopting a hypothesis outside of the 0 to 1 space, we bounce the agent back into the hypothesis space by a distance equaling the overshot amount.

Since our focus here is on problems where the answer is a needle in a haystack, we focus on one particular criterion of epistemic success: whether the best answer to a question (the best treatment, the true causal chain, the correct folding pathway, or the actual treasure location) is found. In terms of our models, that criterion of epistemic success is realized if any individual or group of individuals manages to find the global maximum, rather than some tempting but sub-optimal local maximum. Because all networks at issue are connected, this measure is equivalent to eventual community convergence on the global maximum; if any agent finds the peak, the rest will eventually get there. Given stochastic models such as ours, we can measure how good a particular scientific organization is by the probability that it will find that highest peak.

Our goal here is to understand how the structure of scientific communication and the difficulty of the problem combine to affect scientific outcomes, so in our experiments here we consider a number of different network types in models where the fiendishness index of the epistemic landscape varies.

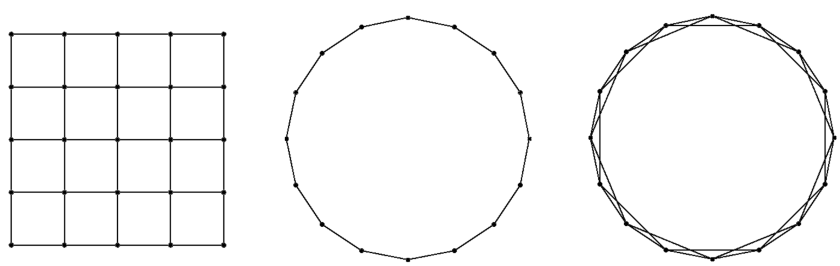

lattice ring radius 1 ring radius 2

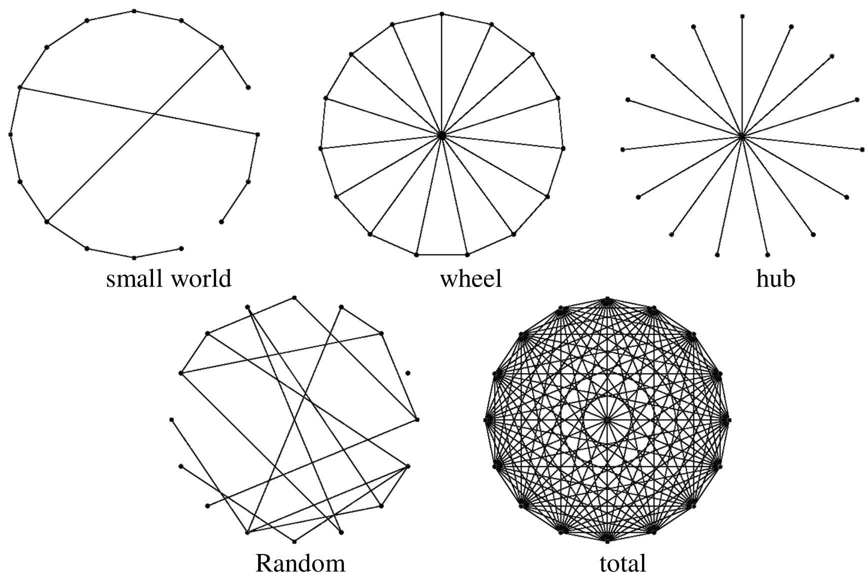

Fig. 6. Types of epistemic network structures. Simulations used a population of 50 nodes (49 in the case of the lattice, which wraps around).

We used 50 nodes (49 in the case of the lattice) and studied the following network structures (illustrated in Fig. 6 with fewer nodes for clarity):

(1) a regular square lattice, with contacts to four agents on each side;
(2) a simple ring, with contacts to a single agent on each side;
(3) a ring with double radius, in which each agent has contacts with two agents on each side;
(4) a small world network, constructed by minimally connecting a simple ring with a 9% probability of rewiring;
(5) a wheel, also known as a 'royal family,' in which each agent on a simple ring also has contact with a central agent;
(6) a hub, in which agents have contact only through a central agent;
(7) a random connected network with a 9% probability of combination between any two nodes;
(8) A total or complete network, in which all nodes are linked to all others.

To explore the effect of network structure on epistemic success, we performed 1000 runs for each network type on landscapes of increasing fiendishness index produced by the formulae given above, with random initial locations of agent hypotheses (and randomized network structure in the case of small world and random networks). For each run, we simulate 100 time steps in which each agent in the network updates its hypothesis based on its feedback using the parameters noted above for speed and margin of error.

Results for these networks on a landscape of the structure of Figure 3 are shown in Fig 7. The x-axis indicates the fiendishness exponent between 0.5 and 10. The y-axis reports epistemic success, that percentage of runs for a particular network in which *any* agent of the network established a hypothesis scoring above 0.85 (i.e., within the region of the global maximum spike above any level possible in the hump).

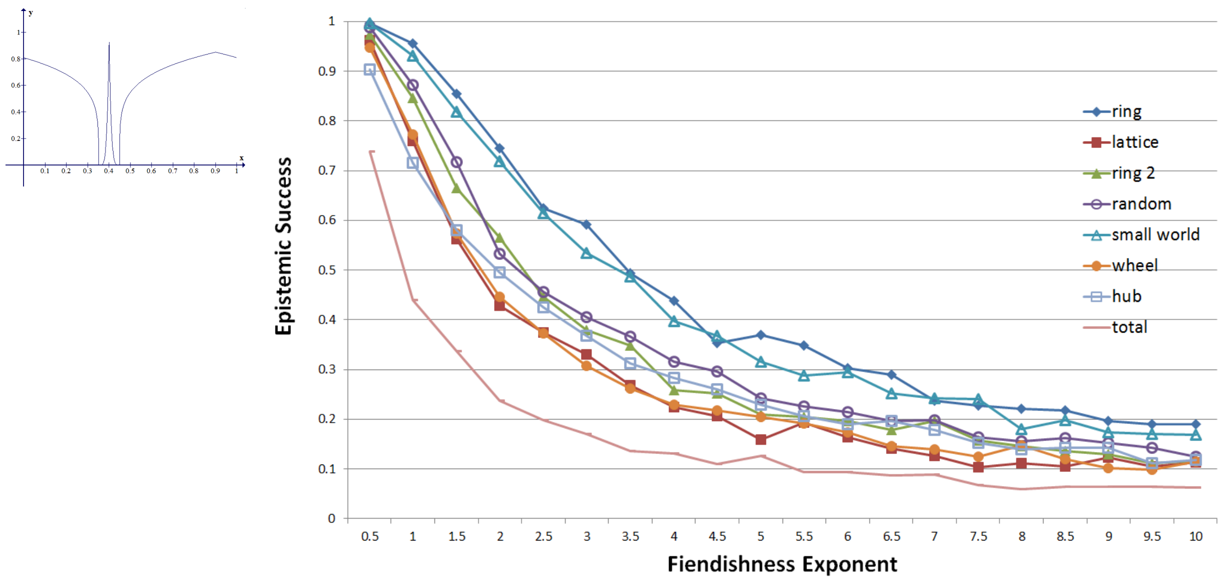

Fig. 7. Epistemic success—percentage of runs in which the spike region is found—across a range of fiendishness exponents and for a variety of networks. Spike center is located at x = 0.4.

Here, the simple ring and small world networks prove most successful across the entire range of landscape fiendishness. Random, hub, double-radius ring, wheel, and lattice networks form a middle group in terms of success, with fully connected networks the clear loser across the entire range. As the problems become more difficult to solve—for high fiendishness—the difference between networks narrows. With a fiendishness exponent of 2, probability differences are dramatic: ring and small world networks have a probability of over 70% of finding the highest peak. The middle group range between a probability of 42% and 58% of finding that peak. Connected networks find the peak in only a little over 20% of cases. At a fiendishness exponent of 3, the ring has a success rate of 60%. The connected network has a success rate of 18%.

It turns out, however, that landscape fiendishness and network structure are not the only factors of importance for this measure of epistemic success, even when assumptions regarding agent updating are held constant. It is not merely the existence of a hidden "needle in a haystack", but also its *position* in the landscape that is of importance to these results.

In two further series we changed the landscape merely by the position of the spike. In one series, shown in Figure 8, the spike was placed precisely in the landscape center (x = 0.5) rather than shifted to one side (at x = 0.4). With that change, the impact of network structure on epistemic success and the relative success of different networks is quite dramatic (Fig. 8).

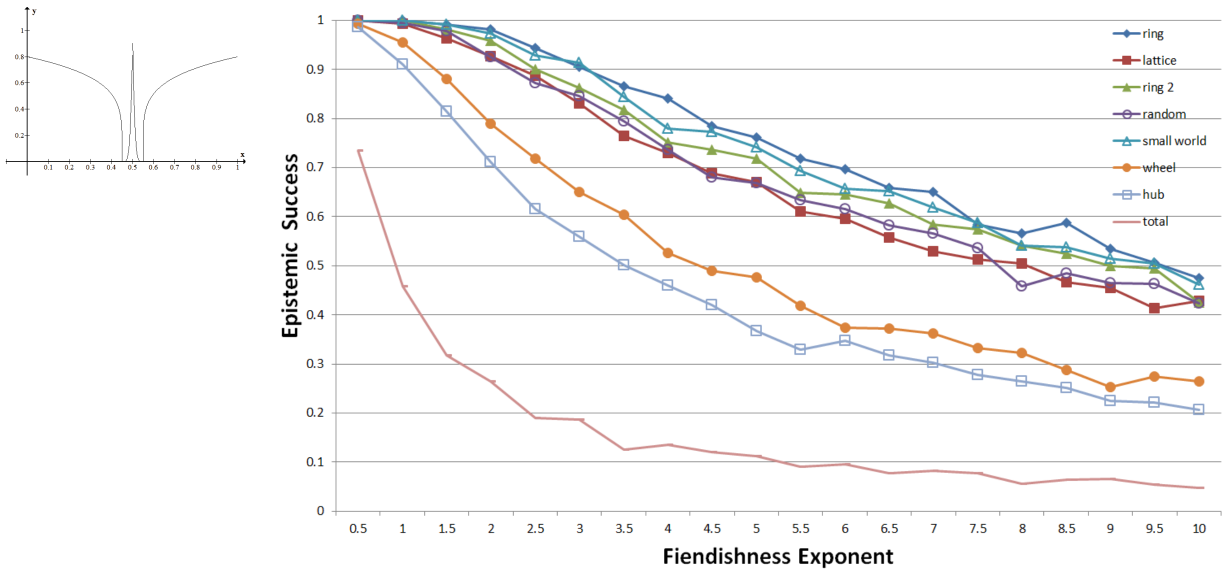

Fig. 8. Epistemic success with spike located precisely in the landscape center.

Fiendishness remains the same in the two cases, since the position of the spike makes no difference to the probability of a highest random drop landing in the spike basin. Given our agent dynamics, however, position of the spike does make a difference in the success of various networks. Ring lattices and small worlds are still more successful than total networks. At a fiendishness exponent of 2, for example, ring and small world networks now have a success rate of above 95%. Total networks can boast a success rate of only a little above 25%. Ring 2 and random networks, which scored significantly below the simple ring in the earlier landscape, here track it fairly closely throughout.

In many cases the order of successful networks has changed as well. In the former landscape, with a shifted spike, the hub scored above the wheel and lattice, which were approximately equal. Here, with a central spike, the lattice is clearly superior to the wheel, which is in turn superior to the hub. A network epistemic success ordering of A > B ≈ C, in other words, has been replaced by C > B > A.

Why the difference? The answer is that our updating dynamics interact with a central landscape structure in a very particular way. With our setting for the speed parameter, approaching the hypothesis of a more successful network neighbor is in steps of 50% toward a target approximated with a margin of error of 10% of the distance. Given a landscape with a centered spike, the peaks of our sub-optimal local maxima fall precisely at 0 and 1 on the scale. If a node is on the sub-optimal hump on one side and its most successful neighbor is higher on the sub-optimal hump on the other side, our dynamics will throw him halfway toward the hypothesis of that more successful neighbor, which will tend to throw him toward the middle of the landscape, increasing the chances of hitting on a more successful hypothesis within the spike.

This result emphasizes the complexity of the phenomenon at issue. Fiendishness clearly plays a role. Network structure clearly plays a role. But precisely what role these play may depend on the interaction of agent updating and other aspects of landscape topography — in particular, the specific position of relevant peaks. This is the particular lesson of dramatic differences in results given the central spike: that not merely the proportional domain but the relative position of epistemic landscape features can spell the difference between success and failure, in some cases reversing the success dominance of particular network structures. Using a very different form of updating, in which networks do not play a pronounced role, Hegselmann and Krause (2006) notes a similarly important role for the "position of the truth" toward the edges or in the center of an opinion space.

We don't think that the positional effect should be regarded as merely a modeling artifact. Epistemic landscapes can vary not merely in proportion of areas of success, but in the relative position of those areas to one another. A successful result may prove easier to find if it is one that is bound to be traversed in the switch from one hypothesis to another. A particular hypothesis may be difficult to hit on not merely because it is a needle in a haystack, but because that needle is located at a particularly unpropitious spot in the haystack as a whole.

The lesson we draw is that there is an additional aspect of landscapes to keep track of, and it is certainly of interest to track the *relative* impact of those different aspects. We also investigated the positional effects by filtering out the effect of the position of the spike by reconstructing our landscape so that it has *no* spatial structure bound between 0 and 1; i.e., we made it "wrap around," thus leaving no center for a spike to occupy. With that change, we're able to extract those network effects that function independently of the positional effect.

The results for a wrap-around landscape of this type are shown in Figure 9, zooming in to focus on the fiendishness exponent from 0.5 to 5.0 in Figure 10.

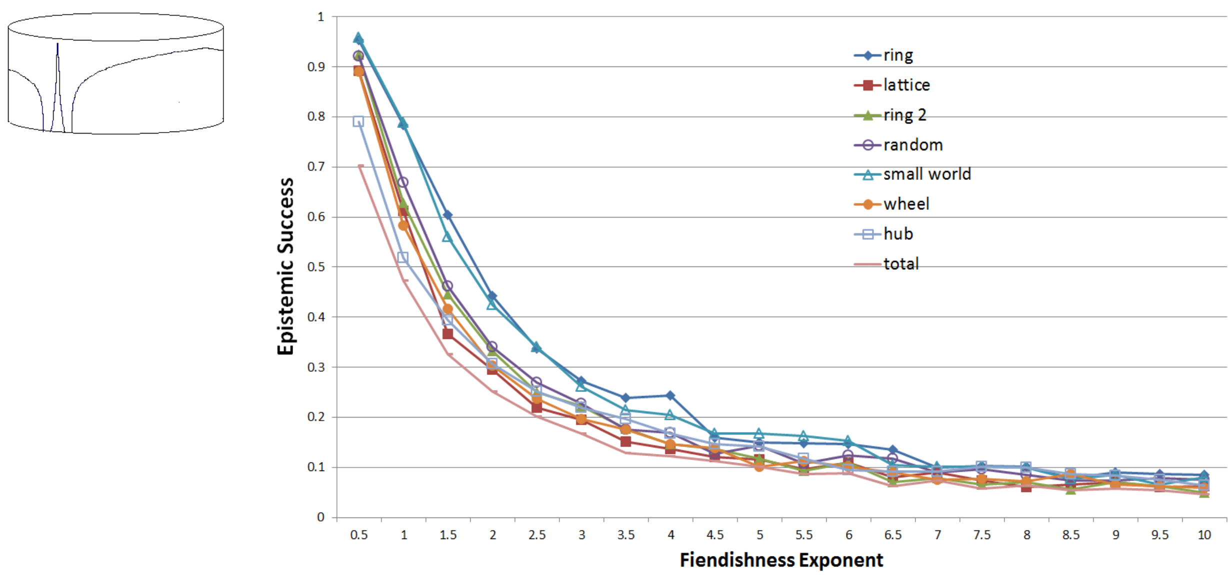

Fig. 9. Epistemic success for a wraparound landscape

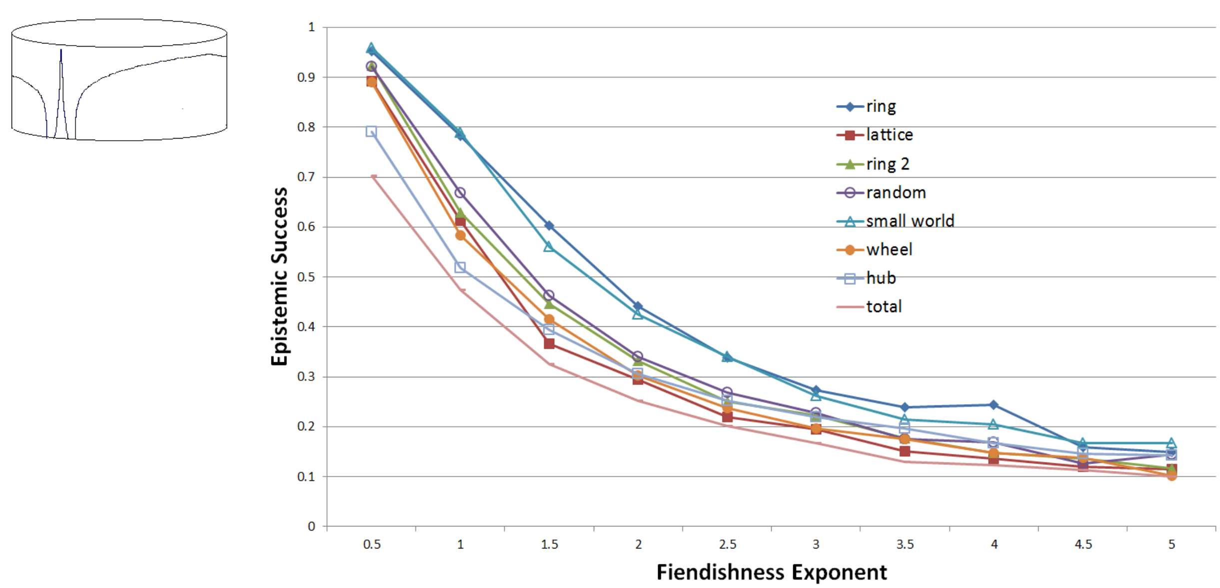

Fig. 10. Epistemic success for a wraparound landscape, focus on fiendishness exponents from 0.5 to 5

Figures 9 and 10 make it clear that results in the other two cases do rely on positional effects. In the case of a wraparound landscape, the divergence of successful from unsuccessful strategies is far less, though at many points the most successful strategies—ring and small world—still have a success rate almost twice that of total networks. At a fiendishness exponent of 1, ring and small world arrays have a success rate of approximately 80%; total networks come in at only 48%. At a fiendishness exponent of 1.5, ring and small world networks approximate a 60% success rate, while total networks are at approximately 32%. In the wraparound landscapes, the order of relative success is the same as our original graph, though the field gets crowded above a fiendish exponent of 4.5 or so.

Interestingly, the case can be made that lattices share the bottom of the pecking order here, precisely as they occupied part of the top region in the case of a central spike. This seems all the more significant since so much work has been done using square lattices (e.g., on cooperation, segregation, and the emergence of language) as if the results were generalizable to networks more broadly (Schelling 1969, 1978; Axelrod 1984; Boccara, Goles, Martinez & Picco 1993; Hegselman & Flache 1998; Cangelosi & Parisi 2002; Grim, Selinger, Braynen, Rosenberger, Au, Louie & Connolly 2005; Grim, Wardach & Beltrami 2006; Grim 2007). The same worry applies to Weisberg and Muldoon (2009) as we mentioned above.

The above results use a particular set of assumptions regarding agent updating to offer a more precise picture of the differential impact of investigatory networks over increasing question difficulty. In doing so, the important and unexpected role of what we have termed 'positional effects' was highlighted.

V. Another Scientific Goal: Convergence

The scientific community puts significant weight on individual discovery — the element of epistemic success we focused on in the last section. But science could not proceed if a discovery remained solely with the discoverer. Dissemination of discoveries allows other scientists to rely on those discoveries and further their own work, and it allows non-scientists to use that information in policy-making and education. The desiderata of both dissemination and discovery are reflected in standard rules governing scientific priority: that the glory goes to the first presentation or publication of a discovery.

The dynamics of the current model guarantee that consensus will emerge since agents always move their hypotheses toward that of more successful neighbors. The amount of time required for convergence of opinion to occur, however, may depend on both the communication network and the epistemic landscape. Kevin Zollman has repeatedly called attention to the trade-off between scientific goals of accuracy and speed (Zollman 2007, 2010a, 2010b, forthcoming). Here we can outline some details of that trade-off against a background of epistemic landscapes with varying degrees of fiendishness.

Recall from above that the ring network is generally the best communication network in finding the hidden spike of a fiendish problem. But as Figures 11 and 12 show, for both the bounded and offset and wraparound landscapes, it is also the ring network in which consensus takes the longest time to achieve.

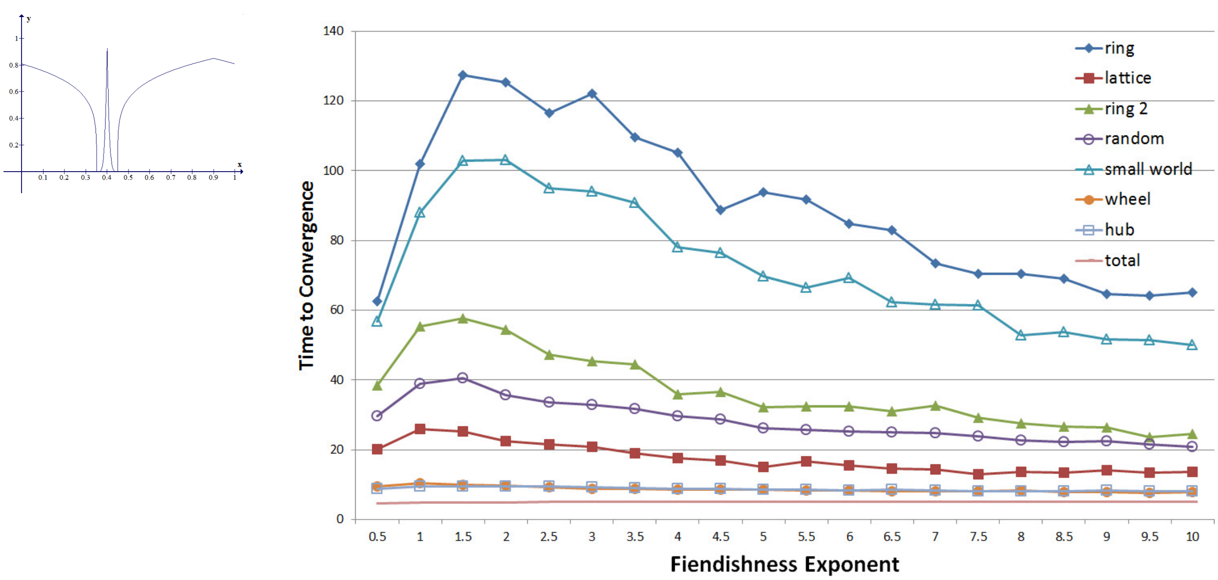

Fig. 11. Time to convergence for a bounded and offset landscape (peak at x = 0.4).

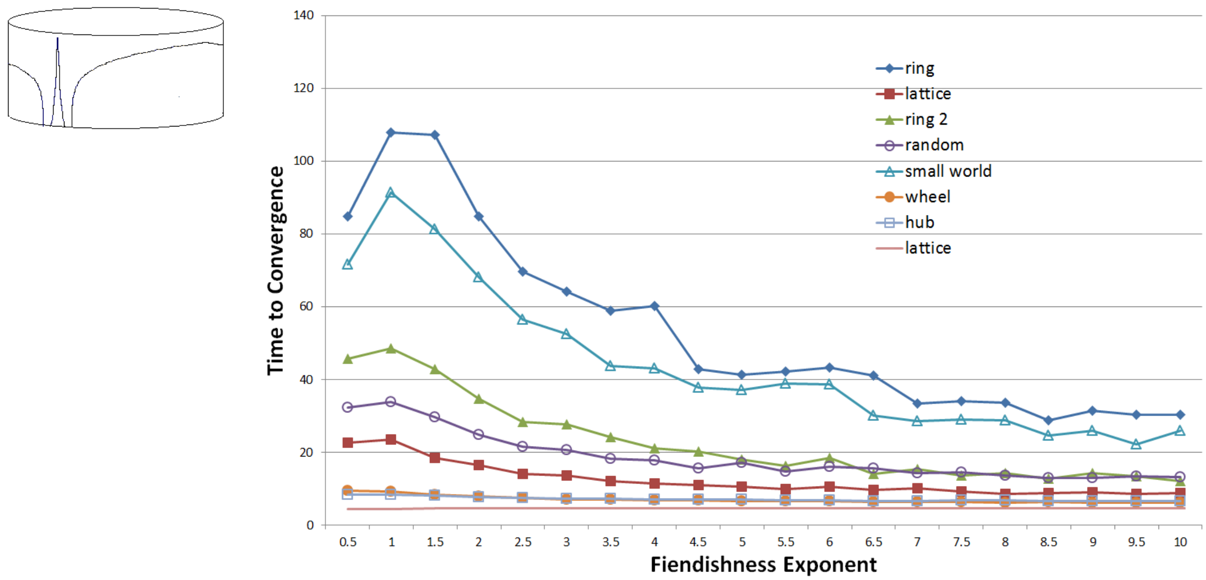

Fig. 12. Time to convergence in a wraparound landscape

If the goal is fast consensus, total networks are clearly optimal. But total networks also do the worst in terms of accuracy. Roughly and qualitatively, the time to consensus and the percentage of runs that find the spike share close to the same ranking: those distributed networks in which exploration is relatively insulated against immediate input from the group are those which have the highest rate of success in individual discovery but also take the longest to achieve a consensus of recognition and adoption.

The precise time to consensus for a given network differs between bounded and offset and wraparound landscapes, just as does the precise percentage of runs in which it finds the highest peak. What holds for both landscapes is the general inverse ordering between accuracy and time to consensus. The higher a network's accuracy on either landscape, the slower its time to consensus. We should therefore expect a bounded landscape with a centered spike, which shows the highest success rates, to show the longest times to network consensus as well. This is indeed what the data shows (Fig. 13).

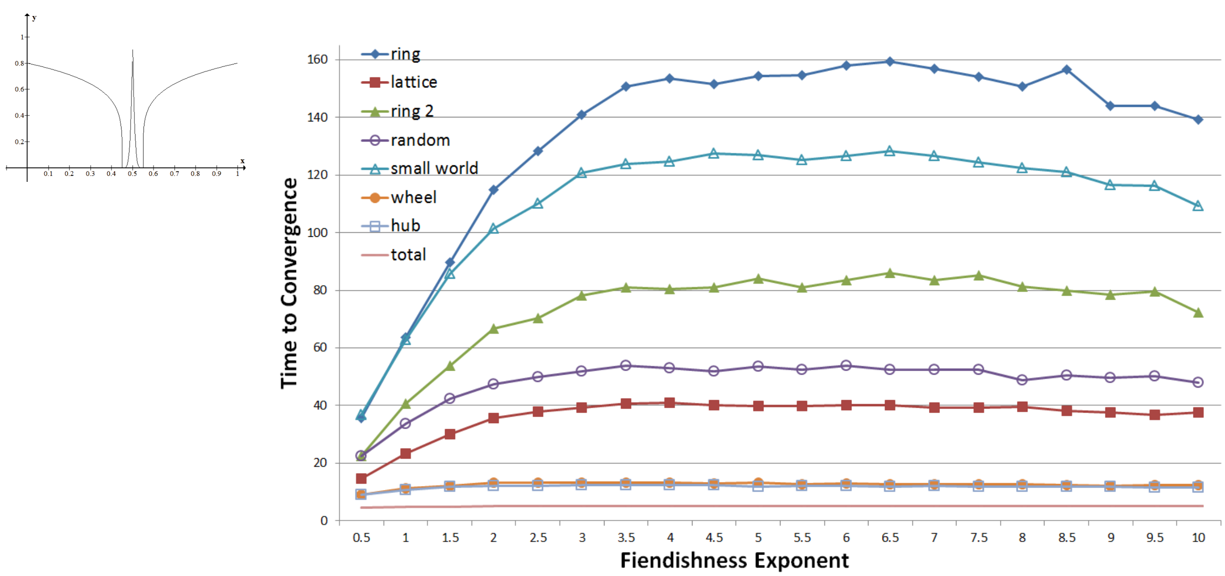

Fig. 13. Time to convergence with a central spike

There is an obvious difference between these graphs. Beyond a fiendishness exponent of 2 or so there is a quick decline in time to convergence in the case of both the wraparound and bounded and offset landscapes, while time to convergence in the case of the central spike shows something closer to a plateau. The declines in Figures 11 and 12 occur because decreasing success, i.e. convergence on the highest point in the landscape, is accompanied by the increasing speed of convergence on the suboptimal hump. If we map instead time to convergence limited to those cases in which convergence was on the highest, we see the appropriate increase in time with increasing question difficulty (Figures 14 through 16). Within those constraints, we see that the convergence results are remarkably uniform.

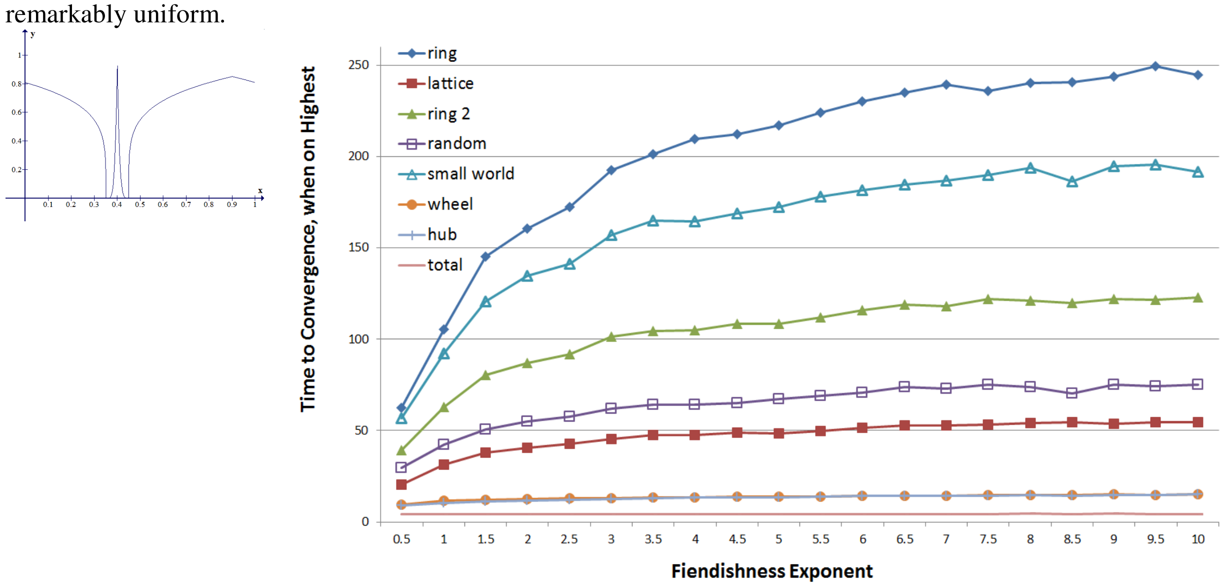

Fig. 14. Time to convergence on a bounded and offset landscape among cases in which convergence is on the highest peak.

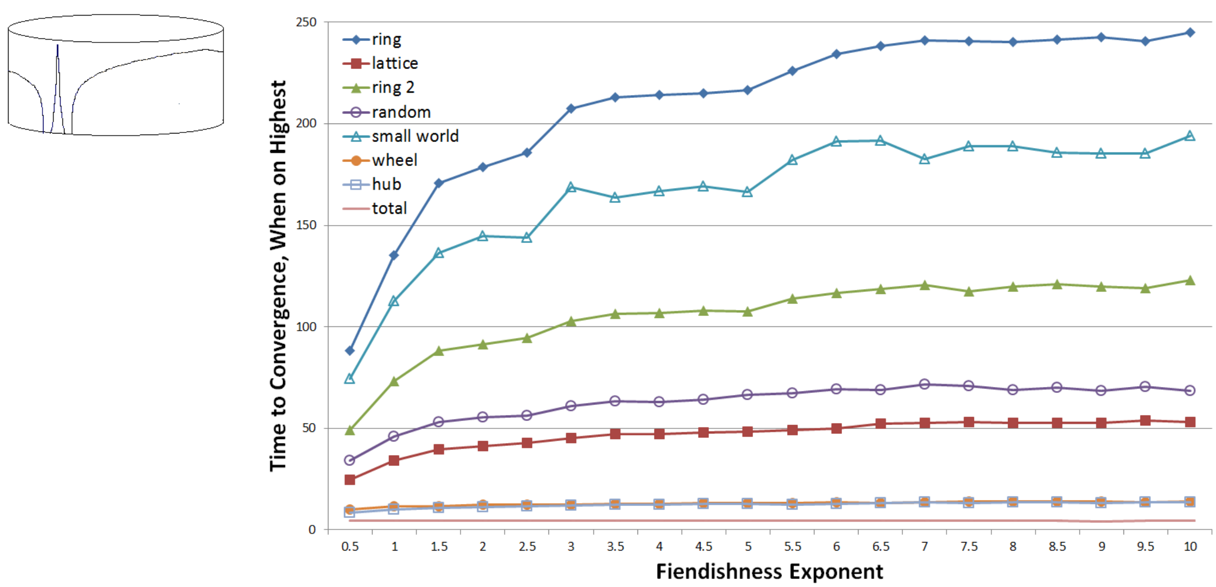

Fig. 15. Time to convergence on a wraparound landscape among cases in which convergence is on the highest peak.

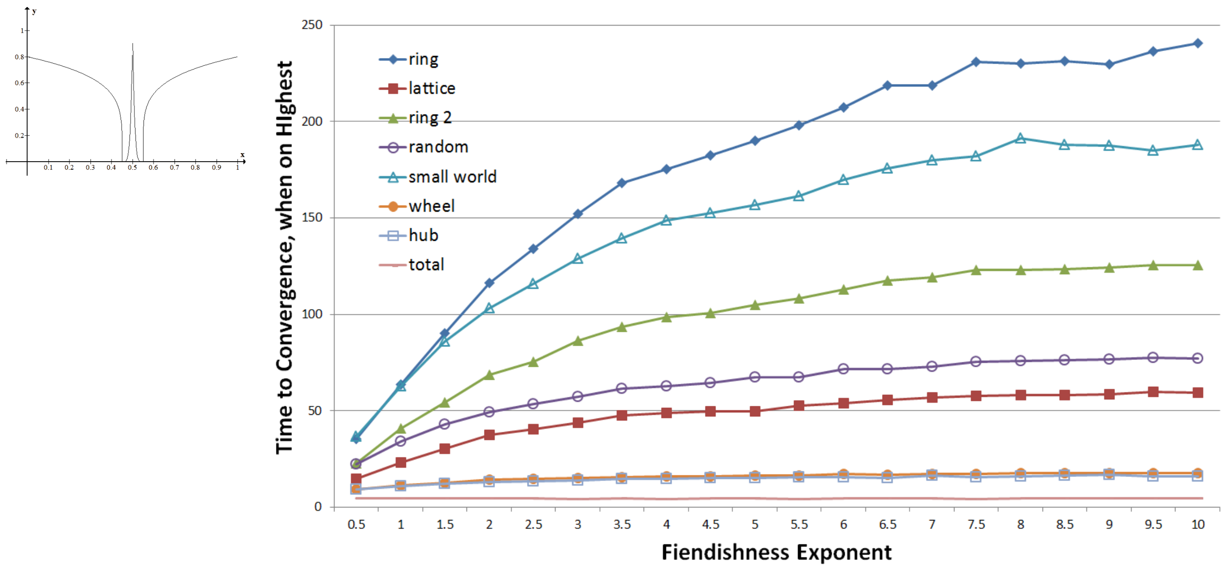

Fig. 16. Time to convergence on a bounded and centered landscape among cases in which convergence is on the highest peak.

VI. Epistemic Success and Convergence on the Watts-Strogatz Continuum

In the previous two sections, we analyzed the impact of network structure on epistemic success (measured in terms of whether any agent finds the peak) and time to convergence. Following Zollman (2007), our method has been to focus on a handful of representative network structures to tease out which aspects of network structure are influence on the target phenomena. Here we introduce a new technique, that of analyzing data produced by a series of models that contain networks produced by systematically varying a network structure parameter.

In the data above, it appears that that median node degree *roughly* tracks the observed results. The highest achieving networks in terms of accuracy are consistently those with the lowest degree; the worst performing is the total, with the highest degree. Those networks which are slowest to achieve consensus are those with the lowest degree; those with high degree converge quickly for precisely that reason.[3]

But neither epistemic success nor convergence are *purely* a matter of degree. Within the limits of keeping rewired networks connected, our ring and small world networks have the same mean degree, track each other closely in terms of success, but diverge importantly in terms of speed to consensus. The 4-lattice and the radius-2 ring networks (in which each node is connected to two neighbors on each side) are both regular networks with a uniform degree of 4. But both success and convergence rates differ significantly between them.

In order to get a better feel for relevant network properties we lay out patterns of epistemic success and convergence in comparison with the well-known Watts-Strogatz model, which produces a series of small-world networks by varying a rewiring parameter. Watts and Strogatz (1998) start with a ring of 1000 nodes, each connected to the 5 nearest on each side. The series of networks is produced by increasing the probability that any link will be rewired randomly, until it reaches 1, where the network is completely random. Across that scale, Watts and Strogatz measure both clustering coefficient — the proportion of node pairs connected to a focal node that are also connected to each other — and characteristic path length — the average number of steps along the shortest paths among all pairs of nodes. When laid out on a log scale, the measure of clustering coefficient across that continuum of increasingly random networks shows an importantly different pattern than does mean path length.

Clustering coefficient and path length are both normalized to their highest value. Watts and Strogatz present their result in terms of a mean of 20 runs. That actually makes the data appear somewhat cleaner than it actually is. The original illustration from Watts and Strogatz appears in Fig. 17.

---

[3] The work of others suggests that degree distribution rather than pure degree can be expected to play an important role in epistemic diffusion, just as it does in patterns of infection (Newman 2002; Meyers, Newman, & Pourbohloul 2006; Bansal, Grenfell & Meyers 2007).

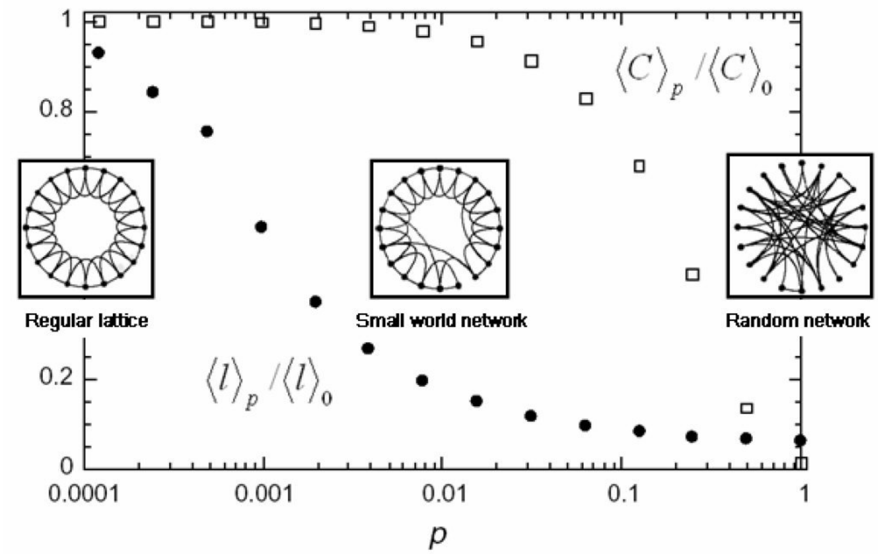

Fig. 17. Normalized clustering coefficient and characteristic path length across increasing probabilities of near-neighbor link rewiring (Watts & Strogatz 1998).

Figure 18 shows our replication of the Watts-Strogatz runs, using probabilities from .0001 and marking values for *each* of the 20 runs. It is clear that mean path length across the spectrum has a significantly wider variance than one might have expected from its original presentation.

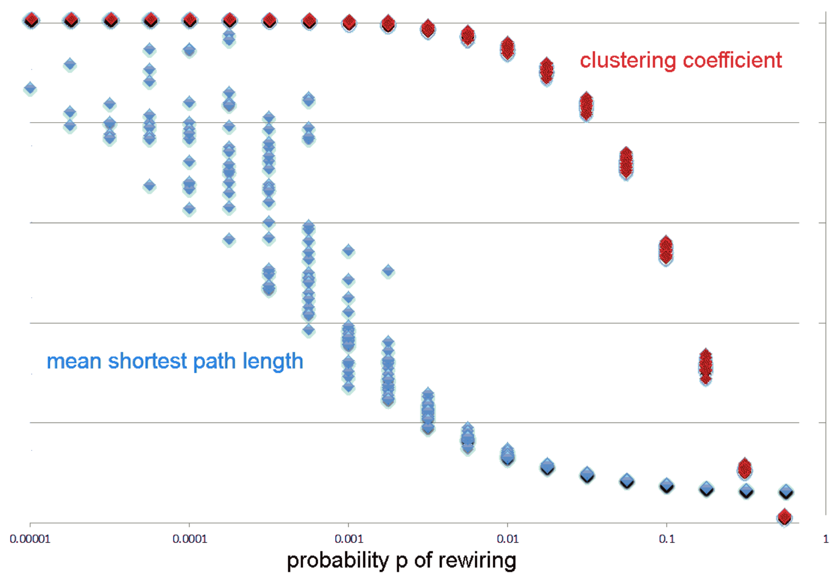

Fig. 18. Replication of the Watts-Strogatz runs, showing the variance in mean shortest path length

Using our replication of the Watts-Strogatz series of networks, we ask whether epistemic success tracks characteristic path length or clustering coefficient. Duplicating Watts and Strogatz's 1000 node networks against a fiendishness exponent of 2.5 for all of our landscapes, it turns out that epistemic success follows neither one.

Across that continuum of networks, the percentage of runs in which the highest peak is found is over 90% in all cases, demonstrating no pattern that suggests that either shortest path length or clustering coefficient closely tracks it. So, for epistemic success, something other than these simple network measures must be at work. Figure 19 shows results for a bounded landscape in which the spike is offset from the center. Figure 20 shows epistemic success for a wraparound landscape. In the case of a bounded and centered landscape, all networks across the continuum had a success rate of 100%. Here the results are shown for the mean of 100 runs in all cases except the bounded and offset network, in which results are shown for the mean of 1000 runs.

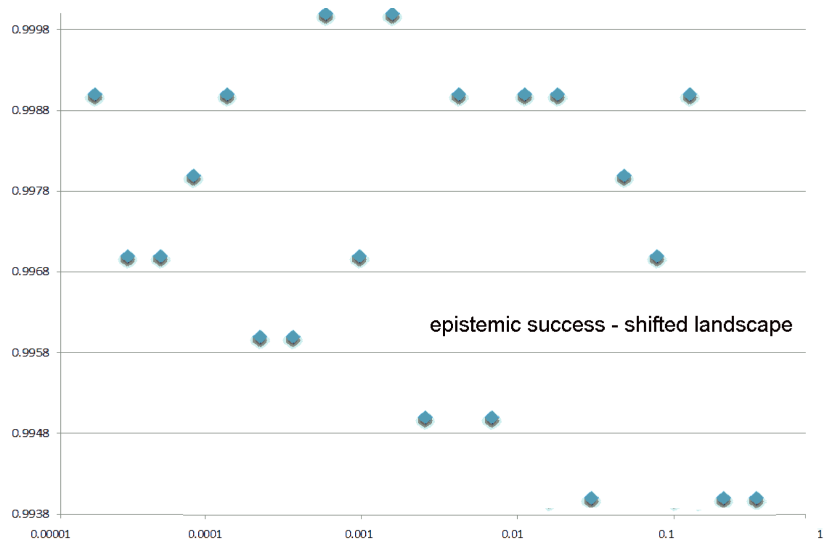

Fig. 19. Epistemic success across the Watts-Strogatz spectrum of networks, using a bounded and offset landscape with a fiendishness exponent of 2.5

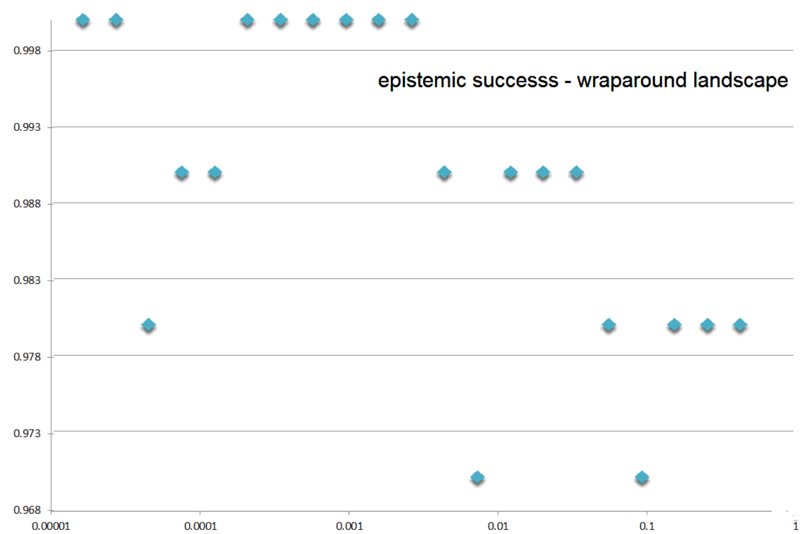

Fig. 20. Epistemic success across the Watts-Strogatz spectrum of networks, using a wraparound landscape with a fiendishness exponent of 2.5

When we turn to the amount of time it takes for the network to converge on a hypothesis, the story is very different. Here the behavior of our epistemic networks *does* follow a simple network property — that of characteristic path length. Figures 21 through 23 show superposition of time to convergence over the characteristic path length obtained in our replication of Watts and Strogatz above.

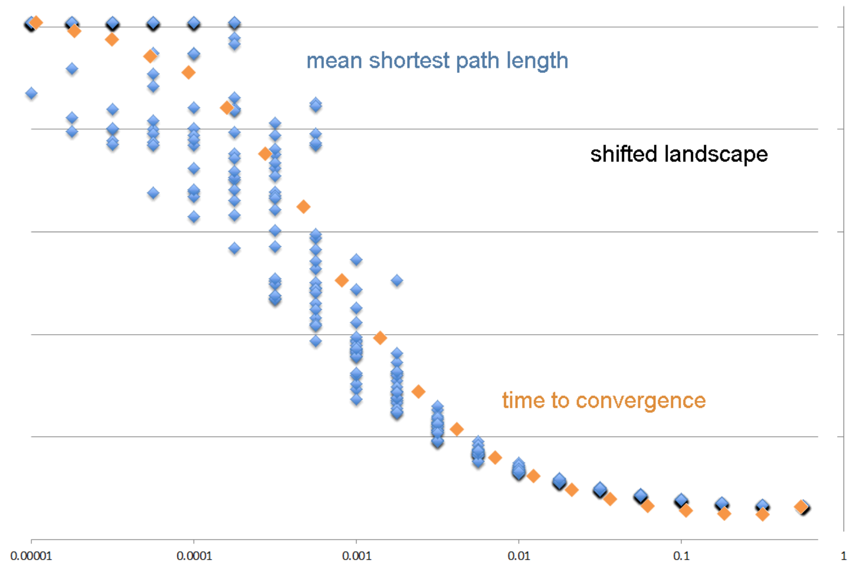

Fig. 21. Time to convergence across the Watts-Strogatz spectrum, using a bounded and offset landscape

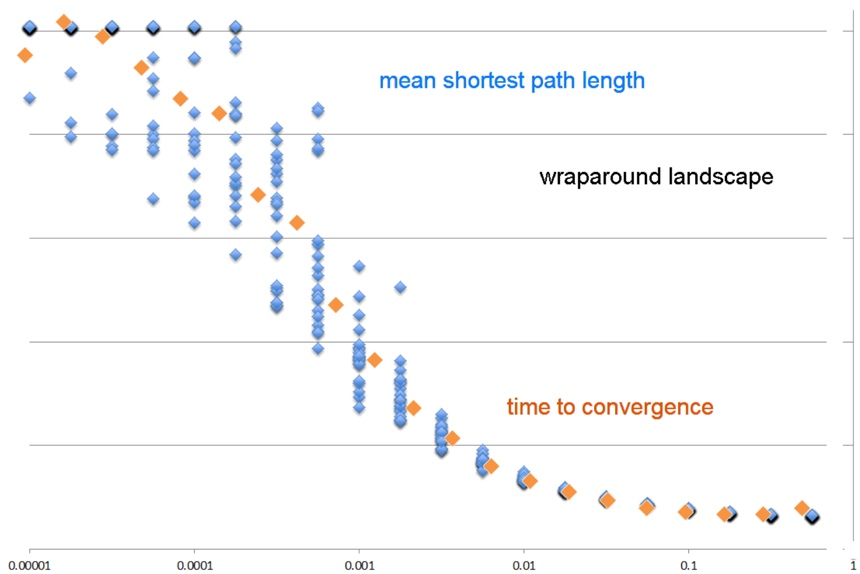

Fig. 22. Time to convergence across the Watts-Strogatz spectrum, using a wraparound landscape

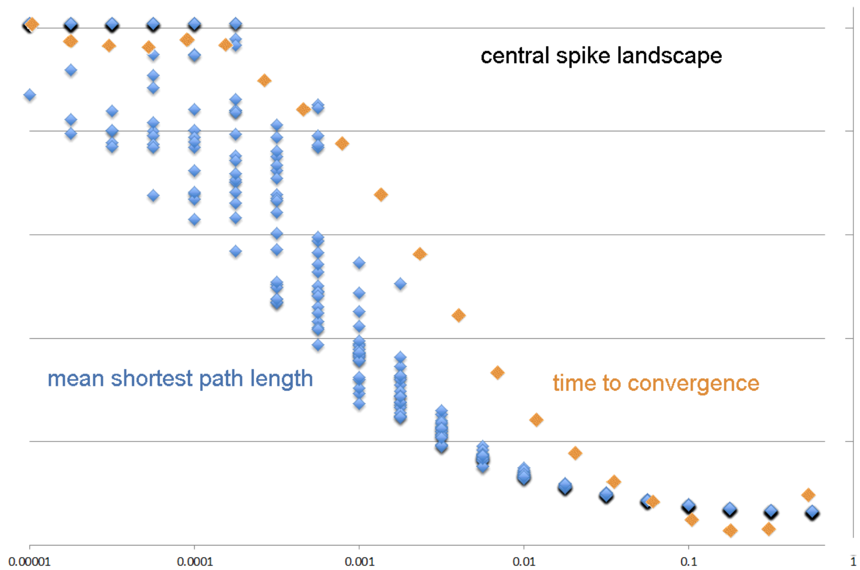

Fig. 23. Time to convergence across the Watts-Strogatz spectrum, using a centered spike landscape

Time to convergence on both the bounded and offset and wraparound landscapes lie squarely along the values for mean shortest path length. Although the curve is the same in a bounded and centered landscape (Fig. 23), the pattern is somewhat different in that case: it is intriguing to note there that time to convergence in that case is slower through the region of small worlds than one would expect from shortest path length alone.

Above, we noted the inverse orderings of successful as opposed to quickly converging networks among our samples, which is consistent with the lesson from Zollman (2007, 2010a). But despite that inverse ordering, the fact that time to consensus tracks characteristic path length in the network, while epistemic success does not, indicates that there are different aspects of network structure at work in the two phenomena. Time to convergence parallels shortest path length quite closely. Epistemic success, even in the simple model we have outlined here, turns out to be significantly more complicated across both our networks and those in the Watts-Strogatz series of networks.

VII. Conclusion

Here we have attempted to shine some light on the complex question of how the social structure of science affects scientific inquiry using models that consist of an epistemic landscape, a social element of communication, and an updating rule that combines inputs from both the landscape and the social element. Previous work using similar models has emphasized the importance of network structure to epistemic success, both in terms of bandit problems (Zollman 2007, 2010a) and difficult epistemic landscapes (Grim 2007, 2009a, 2009b; on a related use of landscapes see Weisberg & Muldoon 2009). Here we have attempted to clarify the relationship more precisely.

In that vein, we introduced the fiendishness index as one measure of question difficulty, more precisely quantifying the "needle in a haystack" quality of a hidden global maxima. In pure terms, the fiendishness index is a measure of the probability that the highest of *N* random points will be within the basin of that highest maxima. But fiendishness alone, it turns out, is not responsible for rates of epistemic success across a sample of networks we investigated. The proportion of area devoted to spike and hump, their curvatures, and their position relative to each other and relative to the boundaries of the hypothesis space can make a dramatic difference in success rates for different networks. For each of three epistemic landscape topologies, we tracked epistemic success and time to convergence properties across increasingly fiendish landscapes. In confirmation of earlier work, these offer a clear trade-off. A full picture of the specific nature of the trade-off and the network properties responsible for it, however, proves harder to find.

To begin to understand what network features influenced success and convergence, we left our handful of sample networks behind and mapped our results onto the Watts-Strogatz continuum of regular to random networks. One of the lessons here, despite the general inverse correlation of epistemic success and speed of consensus, is the important differences between success and correlation across network randomness.

As we mentioned at the beginning, the question of how social structure, individual epistemology, and question difficulty interact in contemporary science is an extremely complex one. To fully understand all of the factors that actually influence science would require resources far beyond what is currently available from psychology, sociology, and the epistemology of science. The simple models here, however, provide a way to begin to understand at least some of the complexities involved.

Suppose we had an indication, or even a rational guess, regarding the fiendishness of a problem. Suppose we had a measure of how important an answer correct in a particular range was for the question at issue. Suppose we demanded consensus, or near consensus, but knew our constraints on time. With inputs of accuracy importance, time constraints, and estimated fiendishness of a problem, could we tell what structure of scientific interaction could be expected to best achieve our scientific and practical goals? We consider all of this a step in that direction.

References

Amaral, L. A. N., A. Diaz-Guilera, A. A. Moreira, A. L. Goldberger, and L. A. Lipsitz (2004). Emergence of complex dynamics in a simple model of signaling networks. *Proceedings of the National Academy of Sciences USA* 101, 15551-15555.

Axelrod, R. *Evolution of Cooperation*. New York: Basic books, 1984.

Bala, V. & S. Goyal (1998). Learning from neighbors. *Review of Economic Studies* 65 (1998), 565-621.

Bansal, S., B. T. Grenfell & L. A. Meyers (2007). When individual behavior matters: Homogeneous and heterogeneous networks in epidemiology. *Journal of the Royal Society Interface* 4, 879-891.

Barrat, A. & M. Weigt (2000). On the properties of small-world networks. *European Physical Journal B* 13, 547-560.

Boccara, N.; Goles, E; Martinez, S.; Picco, P. (1993). Eds., *Cellular Automata and Cooperative Systems*. Nato Science Series, vol. 396. Springer 1993.

Cangelosi, A. & Parisi, D. (2002). Eds, *Simulating the Evolution of Language*. London: Springer.

Conklin, J. (2003). Dialog mapping: Reflections on an industrial strength case study," in P. Kirschner, S.J.B Shum,C.S. Carr Eds, *Visualizing Argumentation – Tools for Collaborative and Educational Sense-Making*, London: Springer-Verlag.

Erdos, P., & A. Renyi (1959). On random graphs. *Publicationes Mathematicae* 6, 290-297.

Erdos, P., & A. Renyi (1960). On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 5, 17-61.

Grim, P. (2006). Tangled webs: The philosophical importance of networks. The Marshall Weinberg Lecture, University of Michigan.

Grim, P. (2007). Network structure in cooperation, communication, and epistemology. Center for Philosophy of Science, University of Pittsburgh, September 2007.

Grim, P. (2009a). Network simulations and their philosophical implications: Models for semantics, pragmatics, and epistemology. Models and Simulations 3, Charlottesville Virginia, March 2009.

Grim, P. (2009b). Threshold phenomena in epistemic networks. *Proceedings, AAAI Fall Symposium on Complex Adaptive Systems and the Threshold Effec*t, FS-09-03, AAAI Press.

Grim, P., E. Selinger, W. Braynen, R. Rosenbergber, R. Au, N. Louie, & J. Connolly (2005). Modeling Prejudice Reduction: Spatialized Game Theory and the Contact Hypothesis. *Public Affairs Quarterly* 19, 95-126.

Grim, P., S. Wardach & V. Beltrami (2006). Location, location, location: The importance of spatialization in modeling cooperation and communication. *Interaction Studies: Social Behavior and Communication in Biological and Artificial Systems* 7, 43-78.

He, J., C. Reeves, C. Witt, & X. Yao (2007). A note on problem difficulty measures in black-box optimization: Classification, realizations and predictability. *Evolutionary Computation* 15, 435-443

Hegselmann, R. "Costs and Benefits of Cognitive Division of Labor and Epistemic Networking: An Agent-Based Simulation Study." Presentation at the Choosing the Future of Science: The Social Organization of Scientific Inquiry Conference. April 20, 2013. Pittsburgh, PA.

Hegselmann, R. & U. Krause (2006). Truth and cognitive division of labour: First steps toward a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation* 9, no. 3. http://jasss.soc.surrey.ac.uki/9/3/10.html

Hegselmann, R. & A. Flache (1998). Understanding complex social dynamics: A plea for cellular automata based modeling. *Journal of Artificial Soceities and Social Simulation* 1**,** <http://jasss.soc.surrey.ac.uk/1/3/1.html>

Kleinberg, J. (2001). Small-world phenomena and the dynamics of information. *Advances in Neural Information Processing Systems* 14, 431-438.

Lazer, D. & A. Friedman (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly* 52: 667-694.

Mason, W. A., A. Jones & R. L. Goldstine (2008). Propagation of innovation in networked groups. *Journal of Experimental Psychology: General* 137, 422-433.

Meyers, L. A., Pourbohloul, B., Newman, M. E. J., & Pourbohloul, B (2006). Predicting epidemics on directed contact networks. *J. Theoretical Biology* 240, 400-418.

Naudts, B. & Kallel, L. (2000). "A comparison of predictive measure of problem difficulty in evolutionary algorithms." *IEEE Transactions on Evolutionary Computation* 4, 1-15.

Newman, M. E. J. (2001). Clustering and preferential attachment in growing networks. *Physical Review E* 64, 025102.

Newman, M. E. J. (2002). The spread of epidemic disease on networks. *Physical Review E* 66, 016 128 (doi: 10.1103/PhysRevE.66.016128)

Rittel, H., & M. Webber (1973). Dilemmas in a general theory of planning," *Policy Sciences*, 4, 155-169. Reprinted in N. Cross (Ed.), *Developments in Design Methodology*, Chichester: J. Wiley & Sons, 1984, pp. 135–144.

Schelling, T. C. (1969). Models of segregation. American Economic Review 59, 488-493.

Schelling, T. C. (1978). *Micromotives and Macrobehavior*. New York: W. W. Norton.

Watts, D. J. & S. H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature* 393, 440-442.

Weisberg, M. & R. Muldoon (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science* 76 (2009) 225-252.

Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science* 74 (5), 574-587.

Zollman, K. J. (2010a). The epistemic benefit of transient diversity. *Erkenntnis* 72 (1), 17-35.

Zollman, K. J. (2010b). Social structure and the effects of conformity. *Synthese* 172 (3), 317-340.

Zollman, K. J. (2012). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, *11*(1), 26-44.