Germs, Genes, and Memes: Function and Fitness Dynamics on Information Networks

Patrick Grim, Daniel J. Singer,
Christopher Reade, and Steven Fisher*†

Understanding the dynamics of information is crucial to many areas of research, both inside and outside of philosophy. Using computer simulations of three kinds of information, germs, genes, and memes, we show that the mechanism of information transfer often swamps network structure in terms of its effects on both the dynamics and the fitness of the information. This insight has both obvious and subtle implications for a number of questions in philosophy, including questions about the nature of information, whether there is genetic information, and how to arrange scientific communities.

Information Transfer Mechanisms: Dynamics and Fitness. Many areas of science are concerned with how information moves on networks. Biologists, for example, try to understand how signals are transferred across animal groups (Handegard et al. 2012; Strandburg-Peshkin et al. 2013) and

Received February 2014; revised August 2014.

*To contact the authors, please write to: Patrick Grim, Department of Philosophy, Stony Brook University, Stony Brook, NY 11794; e-mail: pgrim@notes.cc.sunysb.edu. Daniel J. Singer, Department of Philosophy, University of Pennsylvania, Philadelphia, PA 19104; e-mail: singerd @phil.upenn.edu. Christopher Reade, Gerald R. Ford School of Public Policy, University of Michigan, Ann Arbor, MI 48109. Steven Fisher, Center for Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109.

†This work was supported in part by the National Institute of General Medical Sciences MIDAS grant 1U54GM088491-01, Computational Models of Infectious Disease Threats, administered by the Graduate School for Public Health at the University of Pittsburgh. We are also thankful to many audiences for their feedback on this work, particularly Michael Weisberg's Philosophy of Science Lab at the University of Pennsylvania. Patrick Grim and Daniel J. Singer contributed equally to this research.

Philosophy of Science, 82 (April 2015) pp. 219–243. 0031-8248/2015/8202-0004$10.00
Copyright 2015 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

how genetic information moves and spreads between individuals and across populations in sexual networks (Chowdhury et al. 2010; Holland 2012). Computer scientists ask how to optimally transfer information across computer networks and between networks of sensors (Wooldridge 2002; Germano and de Moura 2006; Tanenbaum 2006; Danon, Arenas, and Díaz-Guilera 2008; Czaplicka, Holyst, and Sloot 2013; *International Journal of Sensors and Sensor Networks*, 2013–). Public health policy makers ask about optimal ways of targeting and delivering health information (Hawe and Potkin 2009; Grim et al. 2010a, 2010b, 2012b). Marketing researchers want to know how information spreads by word of mouth and over social networks and the effect that has on acceptance of innovation and product adoption (Valente 1995; Achrol and Kotler 1999; Gordon 1999; Easley and Kleinberg 2010). Policy makers are interested in information flow across social communication networks with an eye to policy formulation, promulgation, and enforcement (Atkinson and Coleman 1992; Peterson 2003; Nekovee et al. 2007).

These questions also show up in philosophy. How information moves is one of the primary questions of the philosophy of information (Floridi 2003, 2011; Adriaans and van Benthem 2008), but the flow of information is also important to social epistemologists, political philosophers, and researchers on the evolution of norms (Alexander and Skyrms 1999; Goldman 1999; Bendor and Swistak 2001; Bicchieri 2006; Alexander 2007). Additionally, philosophers of science ask how the social structure of science and the transfer of information affect the knowledge of the community at large (Grim 2009; Weisberg and Muldoon 2009; Zollman 2010, 2013; Alexander 2013; Grim et al. 2013).

For concreteness, let us consider three questions:

1. Germs: How does influenza get transmitted between individuals in different social groups?
2. Genes: How does a new advantageous genotype spread through a population over time?
3. Memes: How do communities of scientists converge on a view about the likelihood of a stochastic event when information is sparse?

Each of these questions is about the dynamics of some kind of information on a network. Germs contain genetic information in RNA and DNA that is spread through contact and infection. Genetic information coded in DNA is passed through social reproduction networks. Memes carry propositional information through belief and communication networks.

Because all three questions are about how information moves on networks, they are often lumped together. In particular, it has been popular to draw an analogy between meme transfer and infection dynamics (Le Bon 1897; De Tarde 1903; Park 1904; Park and Burgess 1921; Blumer 1951,

https://doi.org/10.1086/680486 Published online by Cambridge University Press

1969). Douglas Hofstadter, Richard Dawkins, Daniel Dennett, and Susan Blackmore have all discussed beliefs as acting like viruses by spreading socially, competing for fitness, and mutating (Dawkins 1976, 1993; Hofstadter 1983; Dennett 1991; Blackmore 2000).[1] Network and complexity researchers have also adopted this line. Newman, for example, suggests that germs and memes can be considered together by pointing out that varying the structure of the network "might substantially influence the propagation of a rumor, a fashion, a joke, or this year's flu" (2001, 404). Gross, D'Lima, and Blasius also claim that their work on disease spread on networks has implications for "the spreading of information, opinions and beliefs in a population" because memes "can be described in a similar way" (2006, 208701–4).

One way scholars have modeled all three phenomena at once is by treating them all as a simple kind of diffusion on a network. A diffusion model uses a network of nodes that can be thought of as people, communities, bots, journals, or anything else that can serve as bearers of information. The model then progresses step by step: if a node has some information at one time step, then in the next time step its neighbors (probabilistically, at least) have the information. Over time, the information percolates to all of the nodes.

When information flow for germs, genes, and memes is modeled as diffusion, the main determinant of how information moves on the network is the structure of the network. It is well established that network structure does importantly affect the amount of time required for information to spread by diffusion (e.g., Zollman 2007, 2013). What we show in this paper is that any simple model of that kind must leave out important aspects of information dynamics.

Different forms of information have different characteristic mechanisms of transfer. Germs use the asexual reproduction characteristic of infection, genes of higher organisms use sexual reproduction with crossover, and beliefs are transferred through learning, reinforcement, and accommodation. In what follows, we use simple computer-aided simulations to explore the dynamics of these three forms of information transfer on linked subnetworks.

Importantly, our models are idealized and are not meant to be fully accurate models of germs, genes, or memes. As minimal models, they are meant to capture the main dynamical features of the phenomena that give each information transfer type its distinctive character. The purpose of such minimal models is to isolate the core causal features important in explaining the phenomena (Weisberg 2007).

Part 1 shows that transfer type is at least as important as network structure when it comes to convergence speed. Where the timing of the spread across

1. Malcolm Gladwell also makes the comparison in *The Tipping Point*: "Ideas and products and messages and behaviors spread like viruses do" (2000, 7). Lynch (1996), Blackmore (2000), and Brodie (1996) are all book-length elaborations of such a model.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

a community is of concern—whether because we want to delay the spread of viral infection or because we want to optimize scientific inquiry—it is critical that we first ask how the kind of information at issue moves. In part 2 we focus on the relative fitness of these types of information transfer. A highly fit pathogen is one that infects a host population so as to optimize reproduction and spread. A fit gene is one that under selective pressure best survives and spreads across a population. A fit meme or belief is one that accurately represents the world, optimizes cognitive function, or otherwise benefits its agent (Grim 2009). In this second part of the article, we take some first steps toward exploring differences in fitness between these three types of information transfer, again with an emphasis on linked subnetworks. Our results indicate how different types of information can be used to optimize different goals while making clear some of the costs of those optimizations.

Generally, what our results indicate is that questions about information dynamics on networks cannot be answered solely or even primarily in terms of network structure. How network structure affects information dynamics is differentially determined by the specific type of information transfer at issue. We cannot lump together different types of information and ask about their dynamics. It is only once we restrict ourselves to a specific mode of transfer that there are consistent effects of network structure.

Information Networks. Real social networks are not uniform and homogenous. Social communities are composed of subcommunities, with varying degrees of contact in terms of the physical contact necessary for disease transmission, the sexual contact necessary for genetic mixing, and the communication streams necessary for belief transfer. For animals, understanding subcommunities divided by geographical and ecological barriers is crucial for understanding both disease transmission and the genetics of speciation. For people, subcommunities are divided along racial, ethnic, demographic, and socioeconomic lines. In order to understand infection dynamics, we need to understand physical contact networks, including contact links and gaps between subcommunities. In order to understand genetic change, we need to understand links of sexual contact between subcommunities. In the case of belief, we need to understand the impact of linkages between subcommunities, not only of physical contact but also of communication and trust.

Figure 1 shows a series of four networks related in terms of their structure. The network on the left is a single total network. The three pairs on the right form paired total subnetworks with increasing numbers of connecting links.

There are many ways to measure the connectedness of a network, including ratios of the number of links and nodes (like the Alpha, Beta, and

https://doi.org/10.1086/680486 Published online by Cambridge University Press

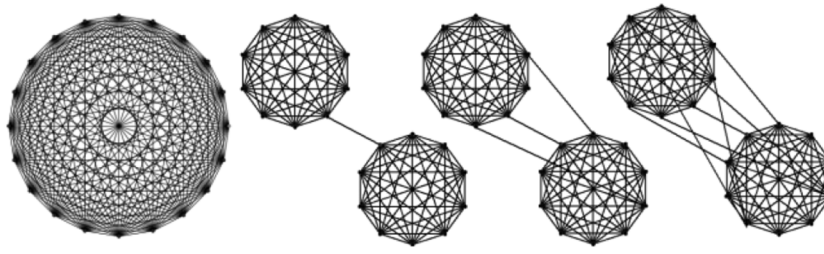


Figure 1. Single total network and increased degrees of linkage between total subnetworks.

Gamma indices); measures over node degrees, betweenness centralities, or closeness centralities of the nodes; graph eccentricity; and average shortest path length. What's relevant to us is the connectedness of the subnetworks to each other, not the connectedness of each subnetwork. As such, we can naturally measure connectedness in terms of the number of linkages between nodes of different subnetworks as a percentage of the total possible. Linkages like this have also been termed 'bridges' in computer networking and identified as a key area for future work in network studies and health care by Trotter, Rothenberg, and Coyle (1995). L. C. Freeman (1978) also uses this notion but describes it as segregation and integration between subnetworks.

Figure 2 shows the types of linked subnetworks we use in the present study: linked total networks, rings, small worlds, and random and scale-free networks. For simplicity, we use just two subnetworks of 50 nodes each. Our rings use just one connection to a single neighbor on each side. For small worlds we work with single rings in which 9% of nodes have been rewired at random. In our random networks 4.5% of possible connections are in each subnetwork. Small-world and random networks had added links to connect them, when necessary. Our scale-free networks are constructed by the preferential attachment algorithm of Barabási and Albert (1999).

In the course of our investigation we vary both the types of subnetworks and the number of links between them.[2] Part of our methodology calls for a comparison between results for linked 50-node subnetworks of a given type and results for a single 100-node network of that type. We compare results for a particular number of links between ring subnetworks with results for

2. To create linked subnetworks, for each bridge, we randomly choose (with repetition) one node from each subnetwork and connect them. This method allows possible 'repeated' bridges, but this shouldn't affect our conclusions since we're comparing linked subnetworks to single networks (which may also have 'repeated' added links).

https://doi.org/10.1086/680486 Published online by Cambridge University Press

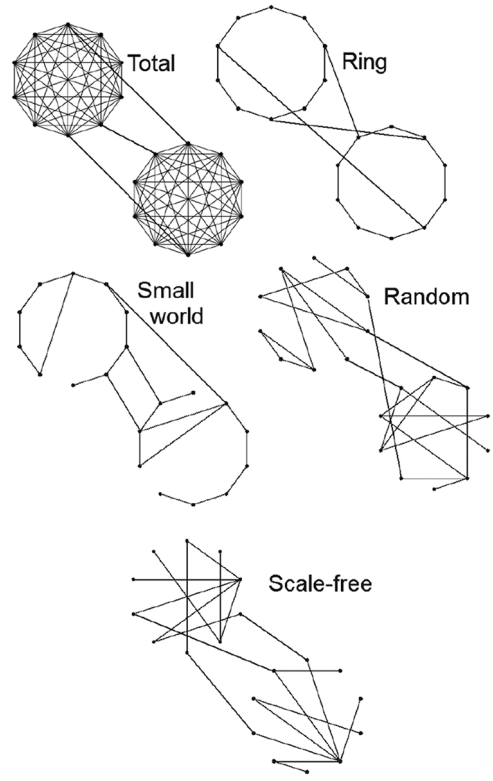


Figure 2. Network types. Smaller numbers of nodes are shown here for visibility.

the same number of links added to a single ring, for example. This method allows us to focus on the effects of two distinct aspects of network structure:

1. network type—ring, wheels, hubs, small worlds, random, or scale-free; and
2. degree of linkage between subnetworks.

We use differences between results on single networks and linked subnetworks to tease out which aspects of overall network structure give rise to particular results. Similar outcomes between a single network of a particular type and linked subnetworks of that type are evidence that it is the network type rather than details of linkage that is important for the result. Different outcomes between a single network and linked subnetworks—particularly where differences carry across different network types—are evidence that it is degree of linkage between the subnetworks that is doing the work.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

Part 1: Network Dynamics for Germs, Genes, and Memes

Infection. Germs are the simplest case. Here we vary (*a*) the structure of subnetworks, that is, whether the subnetworks are rings, small worlds, random, scale-free, or total, and (*b*) the degree of linkage between those subnetworks. What effect does each of these have on the dynamics of infection?

Regardless of the infection transmission rate, time to total infection is importantly sensitive to network structure. It is not sensitive, however, to whether that structure is instantiated as a single network or as linked subnetworks. Figure 3 shows results from simulations for increased linkages between subnetworks. For each number between 1 and 50 we create 1,000 networks with random links of that number between subnetworks, starting with a single infected node. Figure 3 shows the average steps to total network infection over the 1,000 runs, here using a 100% infection rate. Figure 4 shows results in which links are added not between subnetworks but within a single large network of each structure.

Results in the two cases are virtually identical. The difference in plotted lines in each figure shows that network structure does make a significant difference in time to total infection, but the fact that such a structure is instantiated in subnetworks rather than a single network does not make a difference. In all the cases considered, it is network and subnetwork type, not the degree of linkage between subnetworks, that produces the signatures of infection transmission.[3]

To predict, explain, or understand the course of an epidemic—or of information transfer analogous to an epidemic—what one needs to know first is not the degree of linkage between subnetworks but the characteristic structure of the subnetworks themselves.

Memes and Beliefs. Like germs, memes spread across social networks, but in this case, the dynamics of the diffusion are strikingly different. Much work has trumpeted similarities in infection dynamics and the spread of ideas, and they are often treated with the same models (see above and Newman and Watts 1999; Newman 2001; Wang and Chen 2003; Börner, Maru, and Goldstone 2004; Gross et al. 2006; Jackson and Rogers 2007; Pautasso and Jeger 2008). Here we show that they're not as similar as they're often taken to be.

In our models, agents' beliefs are represented as a single number between 0 and 1, which we can think of as an estimate of a fixed quantity, such as about the value of a natural constant, the probability of contracting a disease, or the effectiveness of a vaccination (Janz and Becker 1984; Mullen,

3. This largely accords with analytic results by Golub and Jackson (2012) regarding the role of linkage in diffusion dynamics.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

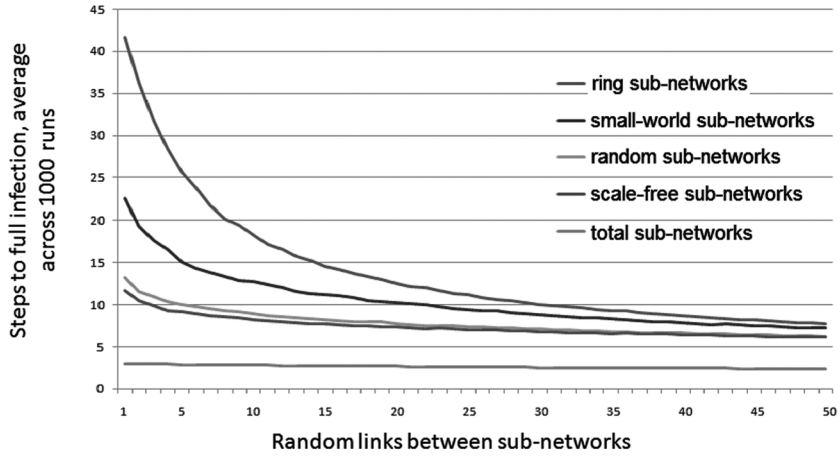


Figure 3. Average time to total infection with increasing links between subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

Hersey, and Iverson 1987; Harrison, Mullen, and Green 1992; Strecher and Rosenstock 1997).

It is clear that a person's beliefs are influenced, among other things, by the beliefs of those with whom they have contact. In our models, agents update their beliefs using an average of the beliefs of those with whom they have

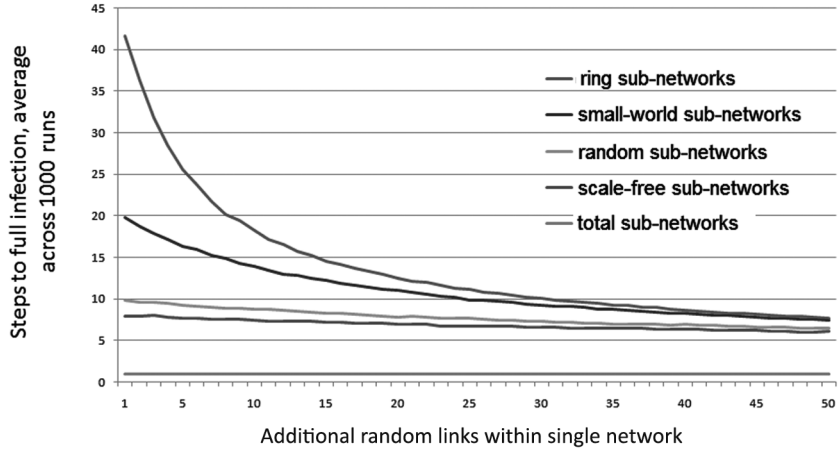


Figure 4. Average time to total infection with increasing links added in single networks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

network contact. Here we follow the central modeling tradition for belief updating first discussed by French (1956) and Harary (1959), brought into a contemporary form by DeGroot (1974) and with network applications, like in Golub and Jackson (2010, 2011) and Dandekar, Goel, and Lee (2013). What an averaging mechanism effectively mimics, in a simple way, is the pattern of reinforcement often characteristic of belief change. The more one's beliefs are like those of more of one's network neighbors, the less inclination there will be to change those beliefs. The more one's beliefs are out of sync with one's neighbors, the greater the pressure there will be to change one's beliefs (Asch 1952, 1955; Bond and Smith 1996; Cialdini and Golstein 2004). DeMarzo, Vayanos, and Zwiebel (2003) argue that updating according to a basic rule of this form can be seen as a boundedly rational heuristic that is consistent with psychological theories of persuasion bias. Alternatives have been proposed, often with attention to parameters of polarization rather than convergence, though these are most easily seen as refinements or variations on the simple form of the central model we use here (Friedkin and Johnsen 1990; Krause 2000; Hegselmann and Krause 2002; Bindel, Kleinberg, and Oren 2011; Acemoglu et al. 2012; Grim et al. 2012a; Su et al. 2014). For simplicity's sake, we concentrate on the flow of beliefs within a communication network, independent of any external input or information, though important work has been done there as well (Zollman 2007, 2013; Grim 2009; Grim et al. 2013).

Using belief averaging, regardless of initial assignment of belief values, all agents in this model eventually approach the same belief value. We can therefore measure the effect of network structure on belief convergence by measuring the number of steps required on average until all agents in the network are within, say, a range of .1 above or below the mean belief across the network as a whole. In what follows we use this range of variance from the mean as our measure of convergence, averaging over 100 runs in each case.

We begin with polarized agents. Half of our agents are drawn from a pool with belief measures that form a normal distribution around .25, with a standard deviation of .06. The other half are drawn from a pool with belief measures in similar normal distribution around .75. In the case of single networks, agents are drawn randomly from each pool. In studying linked subnetworks, our agents in one subnetwork are drawn from the .25 pool; those in the other are drawn from the .75 pool. Belief polarization of this form is necessary to study the effects of subnetwork linkage in particular; were beliefs of our agents randomized within each subnetwork, convergence to an approximate mean could be expected to occur in each subnetwork independently. Time to consensus would not then be a measure of the effect of subnetwork linkage.

In outlining the dynamics of infection above, we contrasted linked subnetworks of particular structures—ring, small world, random, total, and

https://doi.org/10.1086/680486 Published online by Cambridge University Press

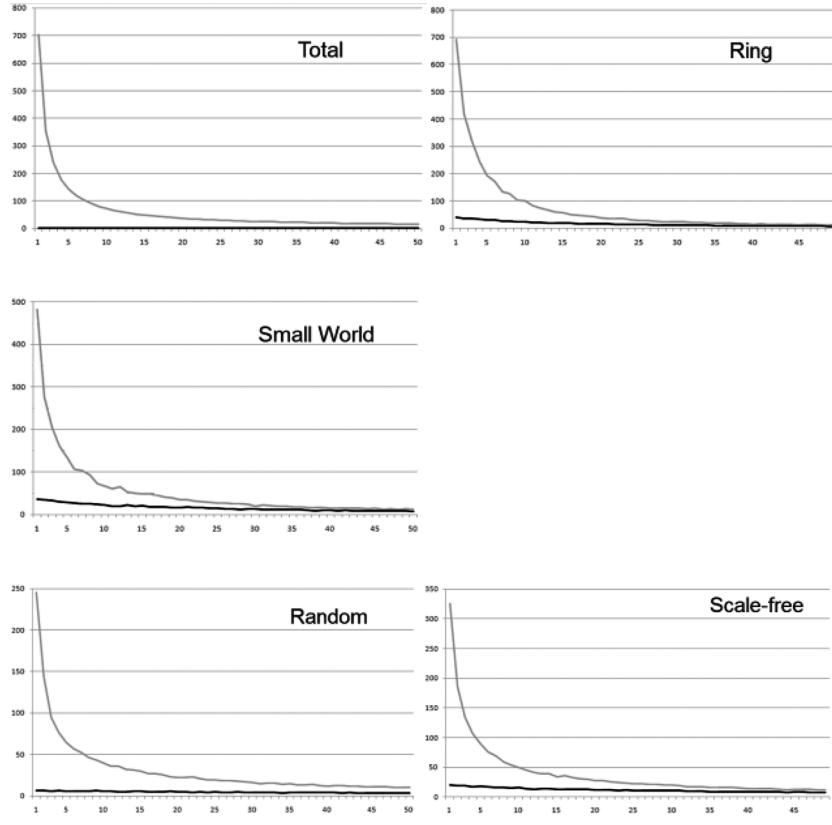


Figure 5. Times to belief convergence in various networks for increasing links between subnetworks (shown in gray) and within single networks of that type (shown in black). Color version available as an online enhancement.

scale-free—with single networks of the same structure. In exploring the dynamics of belief, we again study these side by side. Figure 5 shows graphs indicating times to belief convergence for each of our network types. Times to consensus with increased linkages between subnetworks of a given type are shown in gray. Times to consensus with increased linkages within single networks of that type are shown in black.

Notice two important features of these graphs: (1) the extreme divergence in rates of belief convergence between linked subnetworks and single networks in each case, and (2) the remarkable similarity of the curves for linked subnetworks in each case. That similarity is emphasized by plotting results for all subnetwork types together in log-log form in figure 6, revealing the classic signature of a power law.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

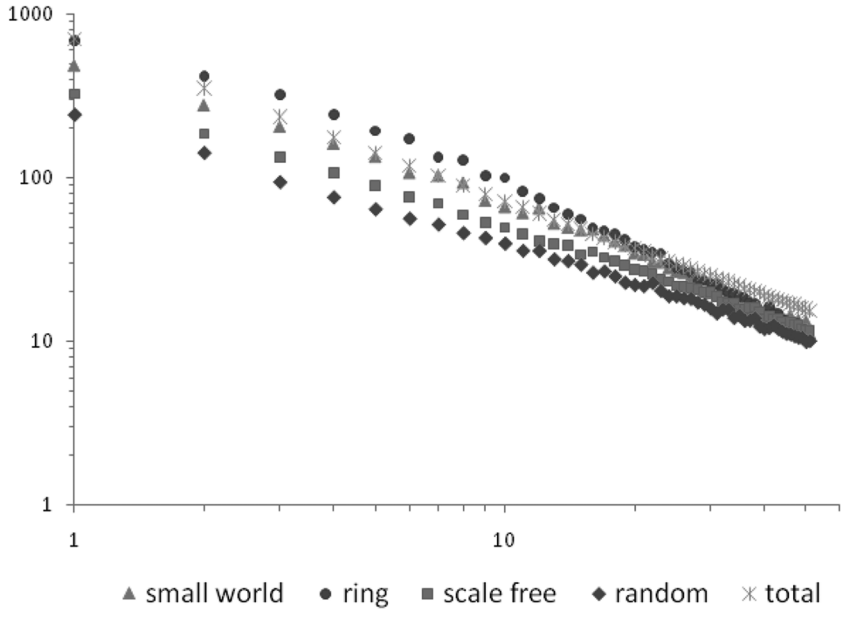


Figure 6. Log-log plots of times to belief consensus with increased linkages between subnetworks. Color version available as an online enhancement.

What these results indicate is that where information is transferred in the manner of memes rather than germs, the dominating effect of increased linkage between subnetworks is independent of the structure of the subnetworks themselves. For meme transmission, unlike infection, the degree of linkage between subnetworks trumps network type. If one wants to trace the course of an epidemic, we noted, it is crucial that one knows the structure of networks involved. To predict, explain, or understand the course of belief transmission, in contrast, it is degree of linkages between subnetworks, of whatever type, that is crucial.

Genes. Genetic information transfer in higher organisms is characterized by crossover in sexual reproduction.[4] To simulate this form of information transfer, we gave each node a genetic code consisting of a binary string of length 100. Half of the population starts out with a genetic code of all ones, the other with all zeros. In the case of linked subnetworks, each sub-

4. Real genetic information and material are transferred in other ways, as well as being mutated. For simplicity, we use 'genetic transfer' to refer solely to the crossover characteristic of sexual reproduction.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

network begins with a uniform genetic code of either zeros or ones. In the case of single networks, we randomize the two codes in the population.

On each time step of the model, each node pairs off with an unpaired node it is connected to, if there is such a node. Each pair then mates, and two new genetic codes are formed. Each new code is a crossover of the two original nodes' codes, code from one node to the left of a random crossover point and code from the other node to the right. The two new codes will generally differ because each code is produced with a random crossover point. Codes for the two original nodes are then replaced with the results of crossover but with the same network connections.

As emphasized in the introduction, our minimal models are meant to contain only the elements that are causally and explanatorily basic, and we don't aim for the models to be perfectly biologically or socially accurate. There are, however, two ways to interpret this model as a model of genetic information transfer. On the first interpretation, the model simulates individual organisms with specific genetic codes mating with those to which they are connected in the network and replaced by their hybridized offspring. One oddity of this way of interpreting the model is that it would have siblings mating sometimes, most notably in the case of ring networks. This oddity is best seen as a consequence of the fact that ring networks are biologically unrealistic. A second interpretation of the model treats each node as a piece of genetic code that is characteristic of a small population of individuals within the larger population. On this interpretation, crossover is interpreted as hybridization of subpopulations, not mating of individual organisms.

In the limit, with connected networks, these models can be expected to converge on a uniform code across the population, which allows us to compare how genetic information moves with how germs and memes move.[5] We measure convergence in terms of the time it takes until two agents drawn randomly from the population can be expected to differ in less than 20% of their genetic code. For our binary strings of length 100, this measure represents a Hamming distance of less than 20.

Above, for belief transfer, we saw that degree of linkage between subnetworks is the major factor in time to convergence. For infection, the crucial factor is network and subnetwork type rather than degree of linkage. Genetic information transfer exhibits a mixture of these features, but it doesn't match either of the other patterns in all respects.

Figure 7 shows genetic dynamics results for linked subnetworks of each of our types. Here, as in the case of belief, network type tends to make very

5. This is not to say that genetic convergence is desirable. Here we simply use the fact that repeated crossover results in convergence to provide a benchmark for comparing genetic information transfer with the other types of information transfer. We don't model the importance of maintaining genetic diversity.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

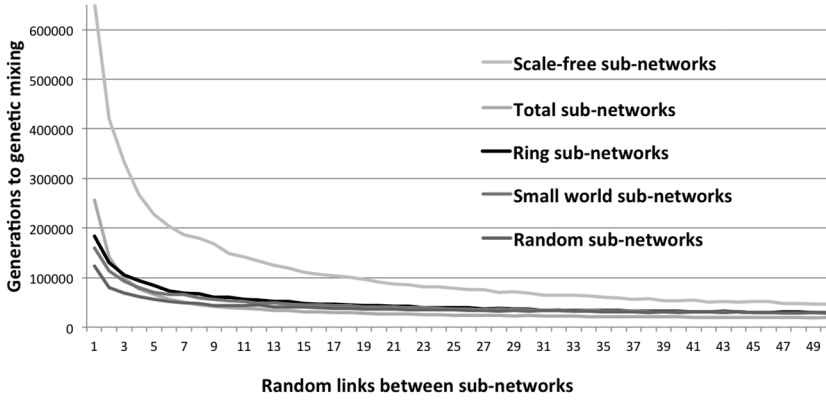


Figure 7. Generations to genetic convergence for increased linkages between subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

little difference. The data from our scale-free preferential attachment networks are outliers; we remain unsure why. When that case is removed, the proximity of results for increased linkages and regardless of the types of subnetworks linked is even clearer. Those results are shown in a rescaled graph in figure 8 and in log-log form in figure 9. Here, as in the case of memes, we have the signature of a power law, though the slope or scaling exponent is very different. Gene transfer therefore shares some notable features with meme transfer.

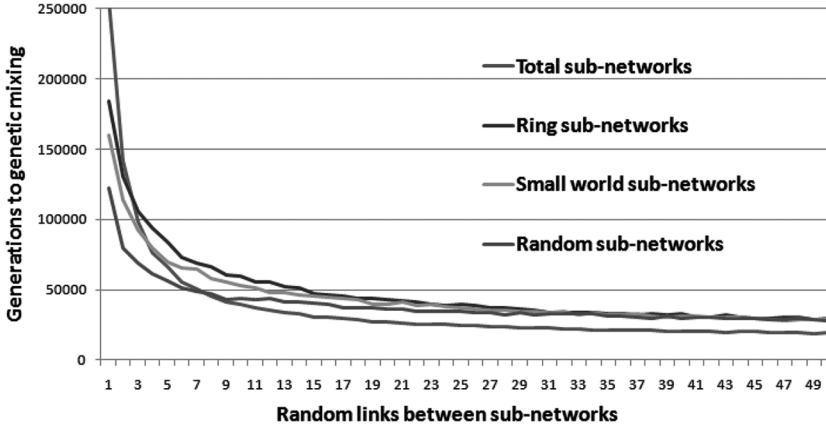


Figure 8. Generations to genetic convergence for increased linkages between subnetworks with scale-free preferential attachment networks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

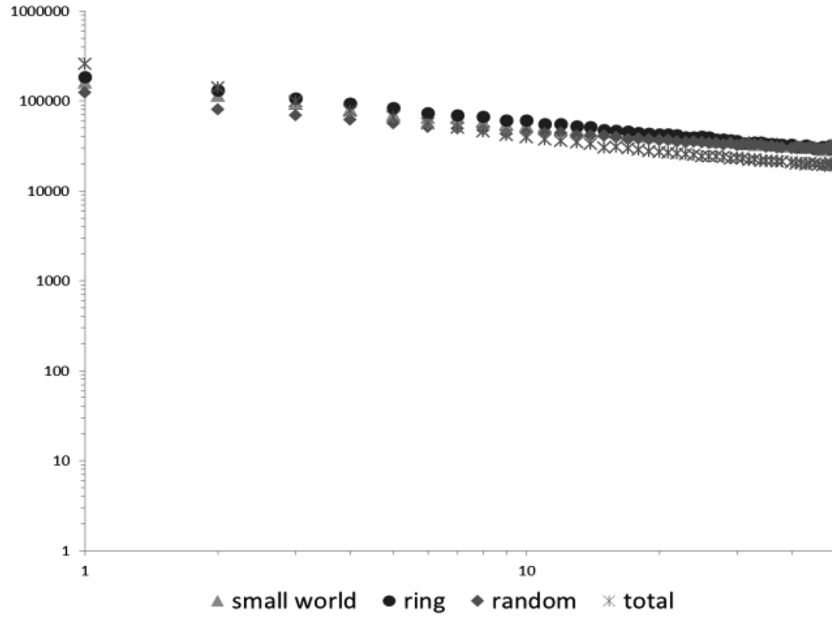


Figure 9. Log-log plot of genetic convergence for increased links between subnetworks. Color version available as an online enhancement.

In the case of memes, however, there is a major difference between results for linked subnetworks and those for single networks of a given type. That difference did not appear in infection dynamics: there, convergence times on linked networks of a given type closely parallel those for single networks with the same number of added links. In this respect, genetic information transfer turns out to be more like infection. Figure 10 shows comparisons for our graph types between added linkages between networks (shown in gray) and within a single network of the same type (shown in black). Total networks stand out, but that is because 'added' links in a total network are redundant. In all other cases single networks start with a lower time to convergence, but after just a few added links, the times to convergence for single and linked subnetworks are nearly identical. In this respect genes are more like germs.

Genetic information transfer therefore shows features of both germ and meme dynamics but is clearly different from both. For memes, the primary determinant of time to convergence is the number of links between subnetworks, not the structure of the linked networks themselves. For germ transfer, the pattern is reversed: network type proves far more important than degree of linkage. For genes, increased network connectivity does in-

https://doi.org/10.1086/680486 Published online by Cambridge University Press

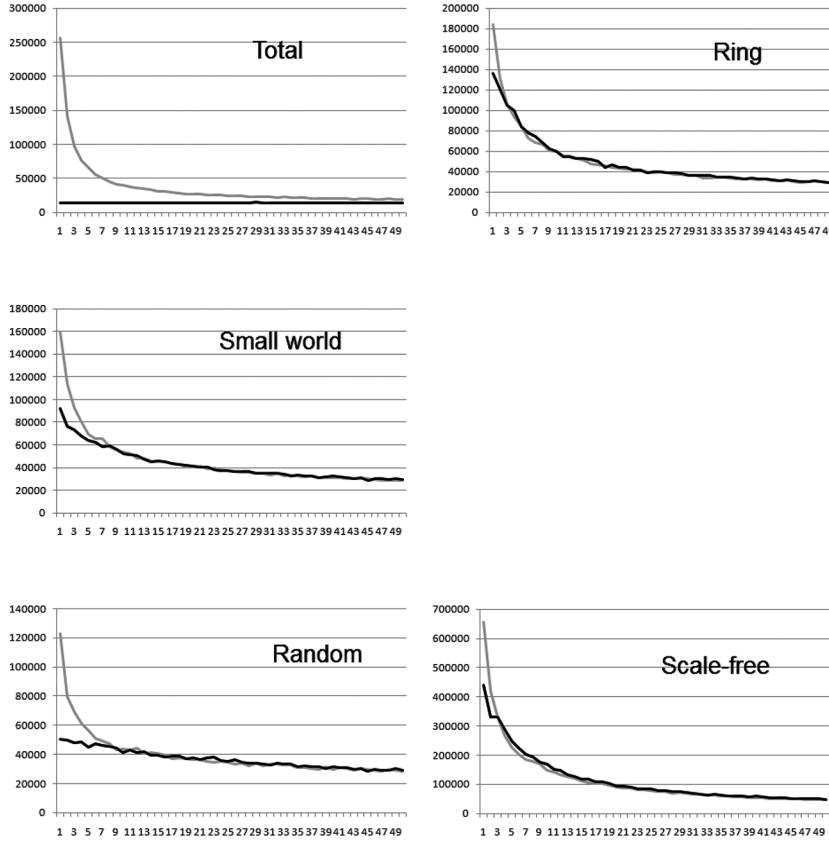


Figure 10. Times to genetic convergence for increasing links between subnetworks (shown in gray) and in single networks (shown in black). Color version available as an online enhancement.

crease speed to convergence, but neither network type nor degree of linkage plays the dominant role characteristic of germs or memes.

Part 2: Fitness for Germs, Genes, and Memes. Different forms of information transfer act in the service of different ends. Here, again using simulations of linked subnetworks as a primary tool, we take some first steps toward a more rigorous understanding of the comparative fitness of different forms of information transfer toward different ends.

We encode information for asexual and sexual reproduction—germs and genes—as binary strings of length 21. In the previous section, we modeled belief as a real number between 0 and 1. By encoding real values for belief

https://doi.org/10.1086/680486 Published online by Cambridge University Press

as binary decimals of length 21, however, we have a consistent way of representing information across all three transfer mechanisms.

Our fitness measure will also be the same across the three mechanisms. We use .10000000000000000000 in binary, or .5 in decimal, as our 'optimal' code: for germs and genes, this can be thought of as the information that is most strongly selected for in a fixed environment.[6] In the case of belief, our optimal code might be thought of as proximity to truth or the ability to guide good action. In all cases, we measure fitness by the arithmetic distance between an agent's code and that arbitrary fitness target.

We use the same range of network types as before, with increased linkages (1) between subnetworks and (2) within single networks of each type. In each case we start with a randomization of information strings across all nodes of the network. What differs is the dynamics of information transfer.

For the asexual reproduction of germs, a node is replaced by that node with the highest fitness to which it is linked in the network. In these early studies we do not include any provision for mutation. For sexual reproduction—genes—the information code of a node is replaced with a crossover at a random point between it and the code of its fittest neighbor, again without mutation. For belief dynamics, the real value of each node is averaged with that of its fittest neighbor.

With a uniform method of representing both information and fitness across the three types of information transfer, we can compare fitness dynamics across different types of networks, whether single or linked subnetworks, and with different degrees of linkage. In each case we measure the number of generations required until the network converges: all nodes have values within .001 of the average fitness of nodes across the network at convergence. At that point we can also measure how 'fit' the final convergence is as proximity to the optimal target at convergence. We use 100 runs for each linkage value of each network, deriving both mean and standard deviation across those results.

The first measure of comparison is time to convergence. In this regard the reproduction of germs is clearly the fastest. Updating by meme or belief reinforcement is the second fastest, with genetic crossover a much slower third. Figure 11 offers just one example: relative speed to convergence with additional links added to a single scale-free network. Slower mechanisms, with more generations to fixation, are displayed toward the top. Similar results hold across all networks in our sample. Germs are faster than memes. Memes are faster than genes.

6. Note that this massively simplifies the notion of fitness in population biology, though treating this optimal string as the genetic code that results in the highest propensity of reproduction, along the lines of the standard account of fitness in population biology, is compatible with our model.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

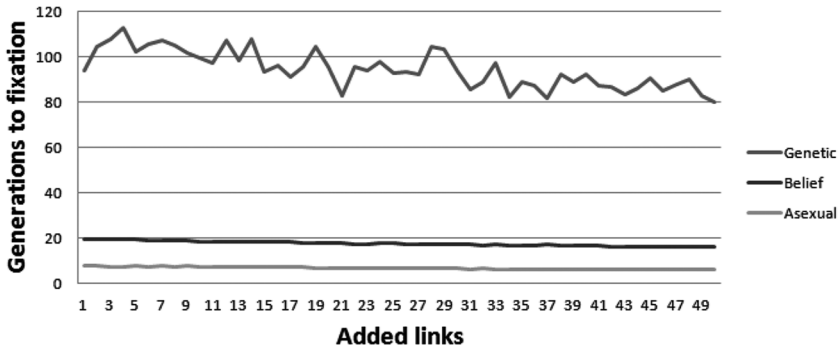


Figure 11. Relative speed to convergence with added links in a single scale-free network, typical of relative speed across all networks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

What of fitness measured as proximity to our optimal code at convergence? Figure 12 shows relative fitness for the three information transfer types, again typical of results for our networks as a whole. The fitter ones, with closer proximity to our optimal value, are displayed toward the top of the graph.

The asexual information transfer of germs is the fastest in terms of time to convergence but scores worst in terms of fitness. The reinforcement dynamics of memes or belief prove the best. The fitness of genetic recombination lies in the middle. On this measure, memes prove fitter than genes, and genes prove fitter than germs.

Orderings in terms of speed and fitness clearly differ for our information transfer mechanisms, but those orderings are not simply inverted. The fast-

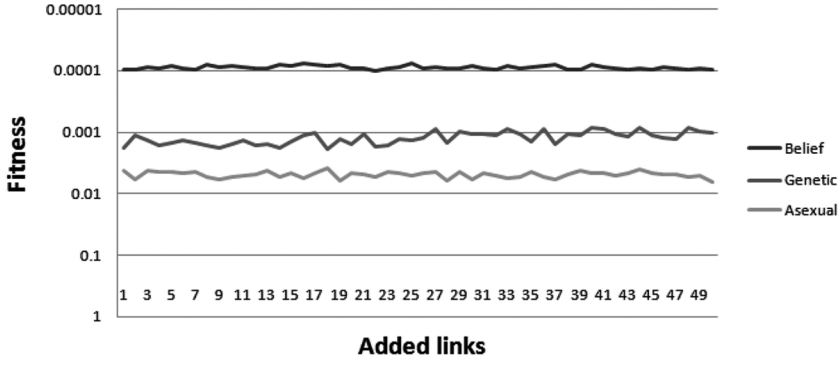


Figure 12. Log-log plot of relative fitness. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

est information transfer mechanism is also that with the lowest fitness: the asexual reproduction of memes. But the slowest transfer mechanism is not that with the highest fitness. The slowest is genetic crossover, but it is the belief reinforcement mechanisms of memes that show the highest fitness.

Across a sample of networks, Zollman (2007, 2010) demonstrates a direct trade-off between speed to convergence and accuracy for networks of agents pooling information about bandit problems. Our work here shows that such a trade-off does not hold across all types of information pooling. Belief transfer proves both faster to convergence and fitter than genetic transfer, for example. This strengthens the lesson from the previous section: to understand the dynamics of information across networks, we must first ask what mechanism of information transfer is under consideration.

There are devils in the details, however, and a more nuanced story is told by the dynamics of different transfer mechanisms on various networks. Figures 13 and 14 show patterns for the asexual reproduction of germs across linked subnetworks of our various types. Speed to convergence varies widely across network type, increasing with increased linkages, without any sharp differences in fitness.

Figures 15 and 16 show comparative results for memes. Linked total networks are the fastest to convergence, but those networks have a very poor fitness. Fitness for other networks is shown in detail in figure 16. With the exception of scale-free networks, for which results are unclear, meme fitness decays with increased linkages. These results are in accord with related work on epistemic networks (Grim 2009; Grim et al. 2010a, 2010b). There, too, increased linkage between subnetworks leads to decay in belief-based information transfer.

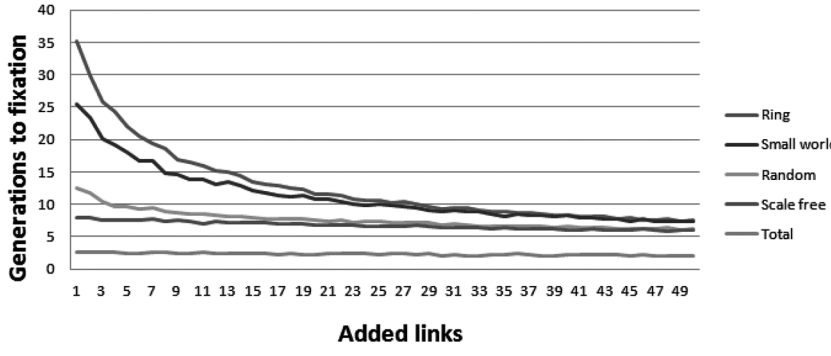


Figure 13. Speed to asexual convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

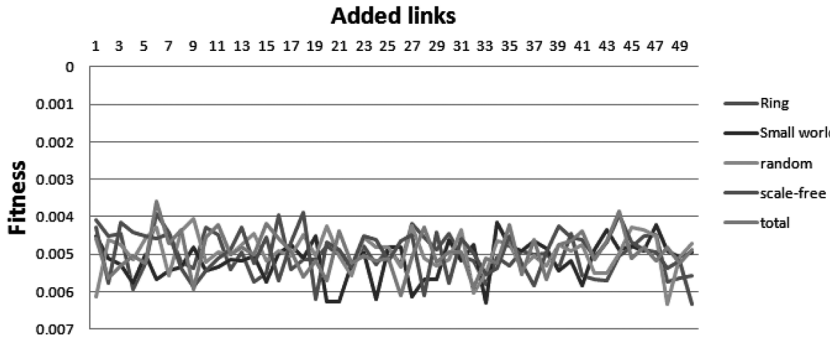


Figure 14. Asexual fitness at convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

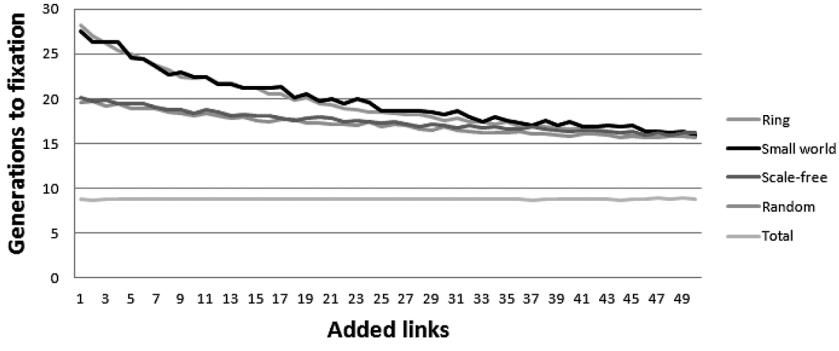


Figure 15. Speed to belief convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

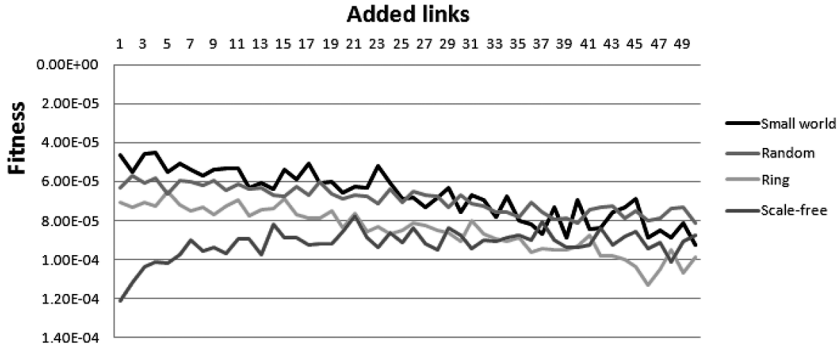


Figure 16. Decay in belief fitness at convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

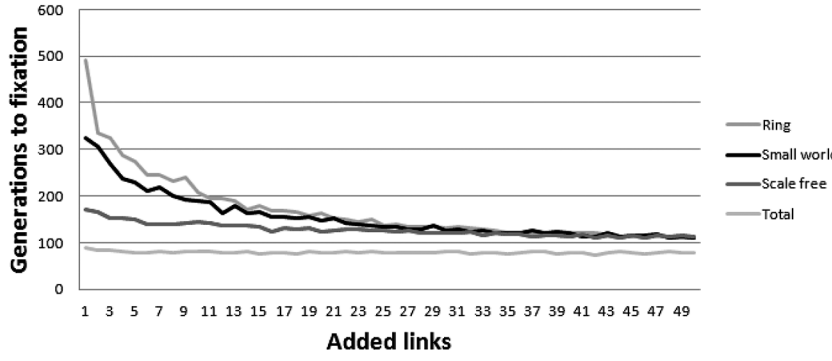


Figure 17. Speed to genetic convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

Comparisons regarding speed and fitness for genes are shown in figures 17 and 18. Though significantly slower, genetic information transfer shows a comparable impact of increased linkages and network types to asexual germs. Like germs, but unlike memes, gene fitness also shows relative constancy across added linkages within a network type. Whereas the fitness of germs is essentially random across network types, however, fitness in the case of genetic transfer does appear to be higher for distributed networks such as rings and small worlds, which echoes a characteristic of memes.

As a whole, our results indicate that information transfer by memes, with the reinforcement characteristic, is an outlier compared to information transfer by either the asexual reproduction of germs or the sexual reproduction of genes. Meme transfer quickly produces convergence and generates highly fit

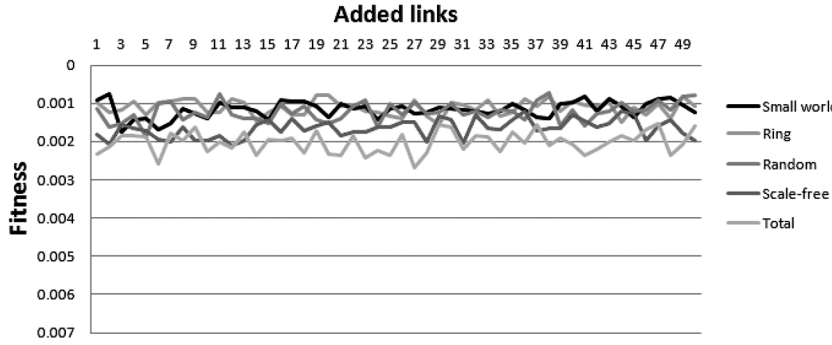


Figure 18. Genetic fitness at convergence with added links between all types of subnetworks. Lines plotted top to bottom, as specified in key. Color version available as an online enhancement.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

communities. But it comes with a cost—a marked decrease in fitness with increased network linkages that does not appear with either germs or genes.

Conclusion. We started out by asking the question, how does information move on networks? Our work here shows that the question cannot be answered independently of a specification of a particular mechanism of information transfer.

The mechanisms of transfer characteristic of germs, genes, and memes have different dynamics in the case of linked subnetworks and different sensitivities to degree of linkage and network type. This means that to understand information flow on networks, we must fix the transfer mechanism before the effect of the network structure is determined. We have taken some first steps in comparing fitness regarding those three types of information as well. Different methods of information transfer optimize different kinds of fitness, selected for in different environments and measured in terms of different ends.

Our results have both obvious and subtle implications for questions in philosophy. Floridi (2011) takes the dynamics of information to be one of the primary aspects of the philosophy of information. One lesson of our work is that there is no single, unified dynamics of information. The picture is significantly more complex: different types of information move importantly differently on different networks.

Our work bears more subtly on questions in social epistemology and the philosophy of science. When setting up institutions for science, we need to understand how propositional information, funding, and credit move on networks of researchers (Kitcher 1990, 1993; Strevens 2003). But the mechanisms of transfer for propositional information, funding, and credit are different. Hence, we should expect that searching for what Zollman calls the "optimal community structure" (2007, 575) may not be best done by merely testing various network structures. We should first understand how propositional information, funding, and credit are each affected by network structure and then combine the results to uncover the optimal social structure for science. Unlike a brute-force search method, this technique allows us to explain why a particular network structure is optimal in terms of the flows of propositional information, funding, and credit across the network.

Finally, in the philosophy of biology, one approach to understanding evolutionary processes and inheritance posits a kind of biological information that is transmitted between generations (Maynard Smith and Szathmáry 1995; Maynard Smith 2000). Others suggest that the supposed 'information' isn't really information at all (Griffiths 2001). Our work can further this discussion and has subtle implications for both sides. We have shown that genetic information transfer does share many formal and dynamic features with memes, a paradigmatic form of information. But in other ways, genetic transfer is also importantly different.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

Because of its importance to questions both inside and outside of philosophy, the issue of information flow has received a great deal of attention. Our attempt here has been to work toward a more cohesive and comparative approach, with a particular point for future research. That point is that one should not expect a unified theory of information dynamics written just in terms of network structure and some single notion of information. Information comes in distinct types with distinct mechanisms. Those differences in mechanism affect both the dynamics and the fitness of information across networks.

REFERENCES

Acemoglu, D., G. Como, F. Fagnani, and A. E. Ozdaglar. 2012. "Opinion Fluctuations and Disagreement in Social Networks." *Mathematics of Operations Research* 38 (1): 1–27.

Achrol, R. S., and P. Kotler. 1999. "Marketing in the Network Economy." *Journal of Marketing* 63: 146–63.

Adriaans, P., and J. F. A. K. van Bentham, eds. 2008. *Handbook of Philosophy of Information*. Amsterdam: Elsevier.

Alexander, J. 2007. *The Structural Evolution of Morality.* Cambridge: Cambridge University Press.

———. 2013. "Preferential Attachment and the Search for Successful Theories." *Philosophy of Science* 80 (5): 769–82.

Alexander, J., and B. Skyrms. 1999. "Bargaining with Neighbors: Is Justice Contagious?" *Journal of Philosophy* 96:588–98.

Asch, S. 1952. *Social Psychology.* Englewood Cliffs, NJ: Prentice Hall.

———. 1955. "Opinions and Social Pressure." *Scientific American* 193:31–35.

Atkinson, M. M., and W. D. Coleman. 1992. "Policy Networks, Policy Communication and the Problems of Governance." *Governance* 5:154–80.

Barabási, A.-L., and R. Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286:509–12.

Bendor, Jonathan, and Piotr Swistak. 2001. "The Evolution of Norms." *American Journal of Sociology* 106:1493–545.

Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.

Bindel, D., J. M. Kleinberg, and S. Oren. 2011. "How Bad Is Forming Your Own Opinion?" In *Proc. 52nd IEEE Symposium on Foundations of Computer Science*, ed. R. Ostrovsky, 57–66. Palm Springs, CA: IEEE.

Blackmore, S. 2000. *The Meme Machine*. Oxford: Oxford University Press.

Blumer, H. 1951. "Collective Behavior." In *Principles of Sociology*, ed. A. M. Lee, 67–121. New York: Barnes and Noble.

———. 1969. *Symbolic Interactionism: Perspective and Method*. Englewood Cliffs, NJ: Prentice Hall.

Bond, R., and P. B. Smith. 1996. "Culture and Conformity: A Meta-analysis of Studies Using Asch's Line Judgment Task." *Psychological Bulletin* 119:111–37.

Börner, K., J. T. Maru, and R. L. Goldstone. 2004. "The Simultaneous Evolution of Author and Paper Networks." *Proceedings of the National Academy of Sciences of the United States of America* 101 (Supp. 1): 5266–73.

Brodie, R. 1996. *Virus of the Mind: The New Science of the Meme.* Seattle: Integral.

Chowdhury, S., J. Lloyd-Price, O.-P. Smolander, W. C. V. Baici, T. R. Hughes, O. Yli-Harja, G. Chua, and A. S. Ribeiro. 2010. "Information Propagation within the Genetic Network of Saccharomyces cerevisiae." *Systems Biology* 4:143.

Cialdini, R. B., and N. J. Goldstein. 2004. "Social Influence: Compliance and Conformity." *Annual Review of Psychology* 55:591–621.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

Czaplicka, A., J. A. Holyst, and P. M. Sloot. 2013. "Noise Enhances Information Transfer in Hierarchical Networks." *Scientific Reports* 3 (1223).
Dandekar, P., A. Goel, and D. T. Lee. 2013. "Biased Assimilation, Homophily, and the Dynamics of Polarization." *Proceedings of National Academy of Sciences* 110 (15): 5791–96.
Danon, L., A. Arenas, and A. Díaz-Guilera. 2008. "Impact of Community Structure on Information Transfer." *Physical Review E* 77 (3): 036103.
Dawkins, R. 1976. *The Selfish Gene*. New York: Oxford University Press.
———. 1993. "Viruses of the Mind." In *Dennett and His Critics: Demystifying Mind*, ed. Bo Dalhbom. Cambridge, MA: Blackwell, 1993.
DeGroot, M. H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association* 69: 118–21.
DeMarzo, P., D. Vayanos, and J. Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics* 118:909–68.
Dennett, D. 1991. *Consciousness Explained*. Boston: Little, Brown.
De Tarde, G. 1903. *The Laws of Imitation*. New York: Holt.
Easley, D., and J. Kleinberg. 2010. *Networks, Crowds and Markets: Reasoning about a Highly Connected World*. Cambridge: Cambridge University Press.
Floridi, L., ed. 2003. *The Blackwell Guide to the Philosophy of Computing and Information*. Oxford: Blackwell.
———. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.
Freeman, L. C. 1978. "Segregation in Social Networks." *Sociological Methods and Research* 6:411–29.
French, J. 1956. "A Formal Theory of Social Power." *Psychological Review* 63:181–94.
Friedkin, N. E., and E. C. Johnsen. 1990. "Social Influence and Opinions." *Journal of Mathematical Sociology* 15:193–206.
Germano, R., and A. P. de Moura. 2006. "Traffic of Particles in Complex Networks." *Physical Review E* 74 (3): 036117.
Gladwell, M. 2000. *The Tipping Point: How Little Things Can Make a Big Difference*. Boston: Little, Brown.
Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford: Clarendon.
Golub, B., and M. O. Jackson. 2010. "Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds." *American Economic Journal: Microeconomics* 2:112–49.
———. 2011. "Network Structure and the Speed of Learning: Measuring Homophily Based on Its Consequences." *Annals of Economics and Statistics*. Available at http://papers.ssrn.com/abstract=1784542.
———. 2012. "How Homophily Affects the Speed of Learning and Best Response Dynamics." *Quarterly Journal of Economics* 127 (3): 1287–338.
Gordon, A. W. 1999. "Network Effects in Marketing." *Marketing Research* 11 (3): 36.
Griffiths, Paul E. 2001. "Genetic Information: A Metaphor in Search of a Theory." *Philosophy of Science* 68 (3): 394–412.
Grim, P. 2009. "Threshold Phenomena in Epistemic Networks." In *Proceedings, AAAI Fall Symposium on Complex Adaptive Systems and the Threshold Effect*. FS-09-03, AAAI Press.
Grim, P., A. Bramson, D. J. Singer, S. Fisher, C. Flocken, and W. Berger. 2012a. "Philosophical Analysis in Modeling Polarization: Notes from a Work in Progress." *American Philosophical Association Newsletter on Philosophy and Computers* 12:7–15.
Grim, P., C. Reade, D. J. Singer, S. Fisher, and S. Majewicz. 2010a. "Robustness across the Structure of Sub-networks: The Contrast between Infection and Information Dynamics." In *Proceedings, AAAI Symposium on Complex Adaptive Systems: Resilience, Robustness and Evolvability.* FS-10-03, AAAI Press.
———. 2010b. "What You Believe Travels Differently: Information and Infection Dynamics across Sub-networks." *Connections* 30:50–63.
Grim, P., D. J. Singer, S. Fisher, A. Bramson, W. J. Berger, C. Reade, C. Flocken, and A. Sales. 2013. "Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence." *Episteme* 10:441–64.
Grim, P., S. B. Thomas, S. Fisher, C. Reade, D. Singer, M. A. Garz, C. S. Fryer, and J. Chatman. 2012b. "Polarization and Belief Dynamics in the Black and White Communities: An Agent-Based Network Model from the Data." *Artificial Life* 13:186–93.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

Gross, T., C. J. D. D'Lima, and B. Blasius. 2006. "Epidemic Dynamics on an Adaptive Network." *Physical Review Letters* 96 (20): 208701.
Handegard, N. O., S. Leblanc, K. Boswell, D. Tjostheim, and I. D. Couzin. 2012. "The Dynamics of Coordinated Group Hunting and Collective Information-Transfer among Schooling Prey." *Current Biology* 22 (13): 1213–17.
Harary, F. 1959. "A Criterion for Unanimity in French's Theory of Social Power." In *Studies in Social Power*, ed. D. Cartwright. Ann Arbor: University of Michigan Press.
Harrison, J. A., P. D. Mullen, and L. W. Green. 1992. "A Meta-analysis of Studies of the Health Belief Model." *Health Education Research* 7:107–16.
Hawe, P., and L. Potkin, 2009. "What Is Population Health Intervention Research?" *Canadian Journal of Public Health* 11 (1): I8–I14.
Hegselmann, R., and U. Krause. 2002. "Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation." *Journal of Artificial Societies and Social Simulation* 5 (3). http://www.jasss.soc.surrey.ac.uk/5/3/2.html.
Hofstadter, Douglas R. 1983. "Artificial Intelligence: Subcognition as Computation." In *The Study of Information: Interdisciplinary Messages*, ed. Fritz Machlup. New York: Wiley.
Holland, J. 2012. *Signals and Boundaries*. Cambridge, MA: MIT Press.
Jackson, M. O., and B. W. Rogers. 2007. "Relating Network Structure to Diffusion Properties through Stochastic Dominance." *BE Journal of Theoretical Economics* 7 (1): 1–16.
Janz, N. K., and M. H. Becker. 1984. "The Health Belief Model: A Decade Later." *Health Education Quarterly* 11 (1): 1–47.
Kitcher, Philip. 1990. "The Division of Cognitive Labor." *Journal of Philosophy* 87 (1): 5–22.
———. 1993. *The Advancement of Science*. New York: Oxford University Press.
Krause, U. 2000. "A Discrete Nonlinear and Non-autonomous Model of Consensus Formation." In *Communications in Difference Equations*, ed. S. Elaydi, G. Ladas, J. Popenda, and J. Rakowski, 227–36. Amsterdam: Gordon & Breach.
Le Bon, G. 1897. *The Crowd: A Study of the Popular Mind*. New York: Macmillan.
Lynch, A. 1996. *Thought Contagion: How Belief Spreads through Society.* New York: Basic.
Maynard Smith, J. 2000. "The Concept of Information in Biology." *Philosophy of Science* 67 (2): 177–94.
Maynard Smith, J., and E. Szathmáry. 1995. *The Major Transitions in Evolution*. Oxford: Oxford University Press.
Mullen, P. D., J. Hersey, and D. C. Iverson. 1987. "Health Behavior Models Compared." *Social Science and Medicine* 24:973–81.
Nekovee, M., Y. Moreno, G. Bianconi, and M. Marsili. 2007. "Theory of Rumour Spreading in Complex Social Networks." *Physica A: Statistical Mechanics and Its Applications* 374 (1): 457–70.
Newman, M. E. 2001. "The Structure of Scientific Collaboration Networks." *Proceedings of the National Academy of Sciences* 98 (2): 404–9.
Newman, M. E., and D. J. Watts. 1999. "Scaling and Percolation in the Small-World Network Model." *Physical Review E* 60 (6): 7332.
Park, Robert. 1904. *Masse und Publikum*. Berlin: Lack & Grunau.
Park, Robert E., and Ernest W. Burgess. 1921. "Collective Behavior." In *Introduction to the Science of Sociology.* Chicago: University of Chicago Press.
Pautasso, M., and M. J. Jeger. 2008. "Epidemic Threshold and Network Structure: The Interplay of Probability of Transmission and of Persistence in Small-Size Directed Networks." *Ecological Complexity* 5 (1): 1–8.
Peterson, J. 2003. "Policy Networks." http://www.ihs.ac.at/publications/pol/pw_90.pdf.
Strandburg-Peshkin, A., et al. 2013. "Visual Sensory Networks and Effective Information Transfer in Animal Groups." *Current Biology* 23 (17): R709–11.
Strecher, V. J., and I. M. Rosenstock. 1997. "The Health Belief Model." In *Health Behavior and Health Education: Theory, Research, and Practice*, ed. K. Glanz, F. M. Lewis, and B. K. Rimer. San Francisco, CA: Jossey-Bass.
Strevens, M. 2003. "The Role of the Priority Rule in Science." *Journal of Philosophy* 100:55–79.

https://doi.org/10.1086/680486 Published online by Cambridge University Press

Su, J., B. Liu, Q. Li, and H. Ma. 2014. "Coevolution of Opinions and Directed Adaptive Networks in a Social Group." *Journal of Artificial Societies and Social Simulation* 17 (2): 4. http://jasss.soc.surrey.ac.uk/17/2/4.html.

Tanenbaum, A. S. 2006. *Computer Networks*. 4th ed. Upper Saddle River, NJ: Pearson.

Trotter, R. T., II, R. B. Rothenberg, and S. Coyle. 1995. "Drug Abuse and HIV Prevention Research: Expanding Paradigms and Network Contributions to Risk Reduction." *Connections* 18 (1): 29–45.

Valente, T. W. 1995. *Network Models of the Diffusion of Innovation*. Cresskill, NJ: Hampton.

Wang, X. F., and G. Chen. 2003. "Complex Networks: Small-World, Scale-Free and Beyond." *Circuits and Systems Magazine, IEEE* 3 (1): 6–20.

Weisberg, M. 2007. "Three Kinds of Idealization." *Journal of Philosophy* 104 (12): 639–59.

Weisberg, M., and R. Muldoon. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76 (2): 225–52.

Wooldridge, M. 2002. *An Introduction to MultiAgent Systems*. New York: Wiley.

Zollman, K. J. S. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5): 574–87.

———. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72 (1): 17–35.

———. 2013. "Network Epistemology: Communication in Epistemic Communities." *Philosophy Compass* 8 (1): 15–27.

https://doi.org/10.1086/680486 Published online by Cambridge University Press