

Information Dynamics Across Subnetworks: Germs, Genes, and Memes

Patrick Grim,^{ad} Daniel J. Singer,^b Christopher Reade,^c & Steven Fisher^d

^aDepartment of Philosophy, Stony Brook University, Stony Brook, NY 11794

^bDepartment of Philosophy, University of Michigan, Ann Arbor, MI 48109

^cGerald R. Ford School of Public Policy, University of Michigan, Ann Arbor, MI 48109

^dCenter for Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109

Abstract

Beyond belief change and meme adoption, both genetics and infection have been spoken of in terms of information transfer. What we examine here, concentrating on the specific case of transfer between sub-networks, are the differences in network dynamics in these cases: the different network dynamics of germs, genes, and memes.

Germs and memes, it turns out, exhibit a very different dynamics across networks. For infection, measured in terms of time to total infection, it is network type rather than degree of linkage between sub-networks that is of primary importance. For belief transfer, measured in terms of time to consensus, it is degree of linkage rather than network type that is crucial.

Genes model each of these other dynamics in part, but match neither in full. For genetics, like belief transfer and unlike infection, network type makes little difference. Like infection and unlike belief, on the other hand, the dynamics of genetic information transfer within single and between linked networks are much the same. In ways both surprising and intriguing, transfer of genetic information seems to be robust across network differences crucial for the other two.

Information on Networks: Germs, Genes, and Memes

Figure 1 shows a series of four network structures, clearly related in terms of structure. The network on the left is a single total network. The three pairs on the right form paired sub-networks with increasing numbers of connecting links. A quantitative measure is possible in terms of the number of actual linkages between nodes of distinct groups or sub-networks over the total possible.

Linkages between sub-networks have also been termed 'bridges,' analogous to a concept of bridges in computer networking and identified in Trotter, Rothenberg and Coyle (1995) as a key area for future work in network

studies and health care. L. C. Freeman (1977) speaks of degree of linkage in terms of segregation and integration between sub-networks.

Our interest focuses on both types of sub-networks and the degrees of linkage between them. What we want to trace is the impact of those aspects of network structure on different mechanisms of information diffusion.




Fig. 1 A single total network and increased degrees of linkage between total sub-networks

In the case of neither people nor animals do realistic social networks form a uniform and homogenous web. Social communities are composed of sub-communities, with varying degrees of contact and isolation between sub-communities in terms of the physical contact necessary for disease transmission, the sexual contact necessary for genetic mixing, and in the case of people the informational contact crucial to transmission of belief.

In the case of animals, sub-communities divided by geographical and ecological barriers are crucial in understanding both disease transmission and the genetics of speciation. In the case of people, sub-communities are also divided along racial, ethnic, demographic, and socio-economic lines. In order to understand belief transfer we need to understand the impact of linkages between sub-communities not only of physical contact but of communication and trust.

Our investigations began with network studies of disease and intervention. Infectious diseases typically exploit social networks; influenza tracks contact networks, while sexually transmitted diseases track sexual networks. Genetic disorders and diseases can be thought of as percolating through genetic networks. Interventions in human disease, on the other hand, typically attempt to

influence behavior change by means of belief change. Health-care behaviors are as crucial in the pattern of any pandemic as are the biological characteristics of the pathogens involved (Epstein, Parker, Cummings & Hammond 2008; Auld 2003; Del Valle, Hethcote, Hyman, & Castillo-Chavez 2005; Barrett, Bisset, Leidig, Marathe, & Marathe 2009; Funk, Gilad, Watkins, & Jansen 2009; Hallett, Gregson, Lewis, Lopman, & Garnett 2007).

Health intervention can thus be seen as a battle on overlapping networks: we attempt to influence the dynamics of disease transmission across contact networks, for example, or genetic drift across genetic networks, by means of information transmission across communication networks. In order to understand prospects for intervention, we need a better understanding of the dynamics of germs, genes, and memes across various types of network structures.

There is also a more abstract way of characterizing our investigations here: in terms of information transfer across networks. Belief change is a clear candidate for information transfer across communication networks, at least if 'information' is drained of any implication of veridicality. But it is also common to speak of genetic information encoded in DNA. We can thus speak of the transfer of genetic information across networks of sexual mixing or hybridization. We can also think of the spread of an infection across a contact network as the spread of information contained in the pathogen at issue.¹ Germs, genes, and memes can thus be seen as exemplifying different forms of information transfer across different kinds of networks. In those abstract terms, our effort is to understand the dynamics of different strategies of information transfer across linked sub-networks.²

A Litter of Linked Sub-Networks

Figure 2 shows the types of linked sub-networks at issue: linked total networks, rings, small worlds, random and scale-free networks. For simplicity we use just two sub-networks of 50 nodes each; Figure 2 uses a smaller number of nodes merely for visibility. Our rings use just one connection to a single neighbor on each side. For small worlds we work with single rings in which roughly 9% of nodes have been re-wired at random. In our random networks roughly 4.5% of possible connections are instantiated in each sub-network. Our scale-free networks are constructed by the preferential attachment algorithm of Barabási and Albert (1999). Where needed, we add a minimal number of links to assure a connected

network in each case; hence the 'roughly' of 9% rewired nodes and 4.5% of possible connections in the case of small world and random sub-networks.³




Fig. 2 The network types at issue. Simulation runs used sub-networks of 50 nodes; smaller numbers of nodes are used here merely for visibility.

We will also be working with single 100-node networks of each type in order to explore the effect of two aspects of network structure:

- (a) network type—ring, wheels, hubs, small worlds, random or scale-free
- (b) degree of linkage between sub-networks.

The idea is to use differences between results on single networks and on linked sub-networks in order to trace those aspects of network structure crucial to particular

¹ Indeed even memes are spoken of as going 'viral.'

² An expanded study would also include diffusion of innovation and marketing strategies across networks (Valente 1995; Goldenberg, Libai & Muller 2001; Garber, Goldenberg, Libai & Muller 2004; Watts & Dodds 2007).

³ Supplementary material regarding hubs and wheels, both analytic and from simulation, can be found at www.pgrm.org/connections

results. Where a result is similar across a single network of a particular type and linked sub-networks of that type, we have evidence that it is the *network type* rather than linkage that is important for that effect. Where results differ in the two—particularly where results are similar between degrees of linkage and regardless of network type—we have evidence that it is the *linkage* between networks that is doing the work.

The Dynamics of Infection

How does the structure of linked networks affect the dynamics of infection across the network as a whole? How does the degree of linkage between networks affect the spread of infection?

It helps to start with analytic results, though we will quickly be forced to simulation.

Consider a pair of total sub-networks connected by a single bridge: two sub-communities that are optimally incestuous internally, but with the slimmest of connections between them. And suppose a single individual in one community becomes infected. Our assumption here and throughout will be of a pathogen with a 100% transmission rate—anyone he touches becomes infected as well. How many steps to total infection across such network?

If our initially infected individual is on the 'outside' of one of the total sub-networks, away from the link, the answer is that infection will spread across the network in just three steps: one to everyone in the sub-network, one across the bridge, and one to everyone on the other side.

If the initially infected individual happens to be on the bridge, saturation occurs in only two steps. The first step simultaneously infects everyone on his side and crosses the bridge to boot. The second step carries the infection from the other side of the bridge to everyone over there.

Where n is the total number of nodes, the average number of steps to total infection in such a case is

$$\frac{3(n-2) + 4}{n} = \frac{(3n-2)}{n}.$$

What if we add more links? All that is going to change is the chance that our 'randomly infected individual' is sitting on the bridge. In all other cases, time to total infection will still be 3 steps.

Adding further links has no dramatic effect in such a case. Because our sub-networks are totally connected, a first step in every case infects all nodes in a sub-network; from there any number of links merely transfer the infection to the second sub-network. For a network with two sub-networks of equal size, incorporating n nodes and

m discrete links between sub-networks (links sharing no nodes),⁴ the average time to total infection will be simply

$$\frac{3(n-2m) + 4m}{n} = \frac{(3n-2m)}{n}.$$

As n increases relative to $m \neq 0$, time to infection approaches a limit of 3. As m increases relative to n , with a limit of $m = .5 n$, time to infection approaches a limit of 2. In all cases, however, variance in infection time will be just between 2 and 3 steps.

We suggested above a comparison with added linkages within a single network. For a total network, of course, 'added' linkages would simply be redundant, with no effect at all. On a single total network, infection at the assumed rate will in all cases be in a single step.

Consider however a very different network structure. What if our sub-communities form not total networks, but rings instead? Infection goes both ways from any starting point. For infection to go to everyone on that ring, we need about 1/2 the number of nodes in that ring. At some point, infection reaches a node on the bridge, which starts the same process on the other side.

Here variance in infection time is much greater. Where s is the number of nodes for a sub-network, the maximal number of steps to full infection from a single node across a ring sub-network is $s/2$ where s is even, and $(s-1)/2$ where s is odd. The longest time for diffusion across a network of two equal-sized rings each with an even number of nodes $n/2$ is therefore

$$\frac{n}{4} + 1 + \frac{n}{4} = \frac{n}{2} + 1.$$

Where the number of nodes $n/2$ in each sub-network is odd the maximal number of steps is

$$\frac{\frac{n}{2}-1}{2} + 1 + \frac{\frac{n}{2}-1}{2} = \frac{n}{2}.$$

If the source of infection is one of the linked nodes, time to infection will be minimal: where $n/2$ is even the minimal time to infection will be $\frac{n}{4} + 1$; where $n/2$ is odd, time to infection will be $\frac{n}{4} + \frac{1}{2}$.⁵

It is clear from even this simple comparison that variance between maximum and minimum times to total infection is extremely sensitive to the structure of sub-networks. In the case of total sub-networks, that variance is simply 1 regardless of the number of nodes. In the case

⁴ In order to keep the outline of basic relationships as simple as possible we ignore the complication that links can share a single node at one end.

⁵ We are grateful to Stephen Majewicz for some of the network analytic used here.

of ring sub-networks, the variance is close to $n/4$. Consequences for prediction are clear: to the extent that a social network approaches a total network, point predictions of infection times can be made with a high degree of confidence. To the extent that a social network approaches a ring, on the other hand, point predictions will not be possible without wide qualification.

Time to total infection is importantly sensitive to network structure. It is not sensitive, however, to whether that structure is instantiated in a single network or in linked sub-networks. Figure 3 shows simulational results for increased linkages between sub-networks. For each number between 1 and 50 we create 1000 networks with random links of that number between sub-networks, taking the average over the 1000 runs.




Fig. 3 Average time to total infection with increasing links between ring, small world, random, scale-free, and total sub-networks

The line for total sub-networks appears at the bottom in Figure 3; in accord with our brief analytical discussion, the result is relatively flat. For linked total networks, average time to infection across the network decreases only 2.98 steps to 2.35 with increased linkages from 1 to 50.

Results for ring networks appear at the top of the graph. For ring sub-networks time to full infection decreases from an average of 38.1 steps for cases in which there is a single link between ring sub-networks to 7.6 for cases in which there are 50 links.

Results for small world networks (with 9% probability of rewiring) are shown in the second line from the top. Increasing linkages from 1 to 50 results in a decrease in steps to total infection from 22.5 steps to 7.45. In the bottom two lines of the graph, virtually the same path is tracked by random networks (using 4.5% of possible links in each sub-network) and scale-free networks (Barabási and Albert preferential attachment). Random networks decrease from 9.79 to 6.45, while scale-free networks decrease from 7.9 to 6.08.

For each network type we ran a parallel series in which additional links are added within a single network of a given structure rather than between sub-networks of that

structure. Results for single networks are shown in Figure 4.




Fig. 4 Average time to total infection with increasing links added to single ring, small world, random, scale-free, and total networks

Figures 3 and 4 are virtually identical. It is clear that network structure does make a significant difference in time to total infection. But the fact that such a structure is instantiated in sub-networks rather than a single network does *not* make a difference. In all the cases considered, it is not degree of linkage between sub-networks but the network structure involved in both single and linked sub-networks that produces network-specific signatures for infection.

This largely accords with analytic results by Golub and Jackson (forthcoming) regarding the role of linkage in diffusion dynamics.⁶ What Golub and Jackson find, working solely with random networks, is that in the limit degree of linkage has no effect in the case of infection or diffusion, propagating by means of shortest paths; in such a case it is only over-all connection density that matters. What our results indicate is that such a result is by no means restricted to random networks, holding across network types quite generally. Where infection is concerned, a prediction of time to total infection demands a knowledge of the general structure of contact network at issue—ring or total, for example, scale-free or random, but does not demand that we know whether it is a single network or a linked set of smaller networks of that same structure that is at issue.

The studies have employed an assumption of 100% infection—a disease guaranteed to be transmitted at every time-point of contact between individuals. More realistic rates of infection of course affect the rates above, but in fact more pointedly emphasize the same points.

⁶ Golub and Jackson characterize their results using the term 'homophily', defined in terms of the relative probability of node connection within as opposed to outside of a group or sub-network. For random networks, though not for other network structures, this corresponds to the degree of linkage between sub-networks that is our focus here.

A more important proviso is that the measure employed throughout has been average time to total infection. Where average time to infection is our measure, degree of linkage between sub-networks as opposed to additional links within a single network of that structure is not of particular significance. But this does not mean that the course of an epidemic across a single network and across sub-networks with various degrees of linkage is not significantly different. That dynamics is often very different—in ways that might be important for intervention, for example—even where average time to total infection is the same. Whereas time to total infection is robust across single and sub-networks, the temporal pattern of that infection is not. The typical graphs in Figure 5 show the rate of new infections over time for (a) a single network and (b) linked sub-networks of that type. Single networks show a smooth normal curve of increasing and declining rates of new infection. Linked sub-networks show a saddle of slower infection between two more rapid peaks.

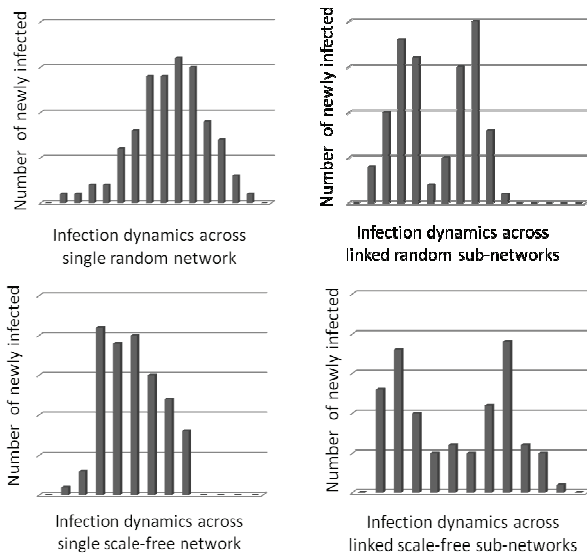


Fig. 5 Contrast in typical dynamics of infection in single and linked sub-networks, even where time to total infection is the same

Despite uniformity of predicted time to total infection, therefore, sparsely linked sub-networks will always be fragile at those links, with temporal saddle points in the course of an epidemic to match. Those weak linkages and saddle points offer crucial opportunities for targeted vaccination in advance of an epidemic, or intervention in the course of it.

The Dynamics of Belief

Like germs, memes spread across social networks. But the dynamics of this form of network information diffusion are dramatically different. Some earlier work has trumpeted similarities in infection dynamics and the spread of ideas (Newman 2001, Redner 1998, Börner et. al. 2003). Our purpose is to emphasize crucial differences between them.

In this first model our agents' beliefs are represented as a single number between 0 and 1. These are beliefs in the severity of a disease, perhaps, the probability of contracting the disease, or the effectiveness of vaccination. (Harrison, Mullen, & Green 1992; Janz & Becker, 1984; Mullen, Hersey, and Iverson, 1987; Strecher & Rosenstock, 1997). Agents are influenced by the beliefs of those around them, updating their belief representation in terms of the beliefs of those with whom they are informationally linked.

To this extent we can argue that the model is relatively realistic: some beliefs can be represented on such a scale, and people are influenced to change those beliefs by, among other things, the expressed beliefs of those with whom they have contact. What is admittedly unrealistic is the simple form of belief updating we use in the model: an averaging of current beliefs with those with whom one has network contact. No-one thinks that averaging of beliefs in an informational neighborhood captures the real dynamics of belief change. Such a mechanism does, however, instantiate a pattern of reinforcement: the more one's beliefs are like those of one's network neighbors, and the more they are like 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. Our attempt is not to reproduce any particular pattern of realistic belief change but to emphasize the impact of certain predictable characteristics of belief change—with reinforcement a primary component—on the dynamics of belief change across communication networks.⁷

Given belief averaging, and regardless of initial assignment of belief representations, 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,

⁷ Centola and May (2007) consider 'complex contagions', in which more than one neighbor is required for infection. This is not strictly speaking a reinforcement effect, but does show dynamics similar to that studied for belief reinforcement here—and a similar contrast with simple infection. Golub and Jackson (forthcoming) outline analytic results on homophily in random networks, with a similar contrast between diffusion and belief averaging. Our results extend that work, indicating that the central contrast holds across networks of various types.

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 deviation of .06. The other half are drawn from a pool with belief measures in similar normal distribution around .75. In studying linked sub-networks our agents in one sub-network are drawn from the .25 pool; those in the other are drawn from the .75 pool. In the case of single networks agents are drawn randomly from each pool. We found belief polarization of this form to be necessary in order to study the effects of sub-network linkage in particular; were beliefs of all our agents merely randomized, convergence to an approximate mean could be expected to occur in each sub-network independently, and time to consensus would not then be an adequate measure of the effect of sub-network linkage.

In outlining the dynamics of infection we contrasted linked sub-networks of particular structures—ring, small world, random, total, and scale-free—with single networks of the same structure. In exploring the dynamics of belief we will again study these types side by side.

Figure 6 shows graphs with indicating times to belief convergence for each of our network types, for increased linkages between sub-networks of that type (shown in blue) and increased linkages within single networks of that type (shown in red).

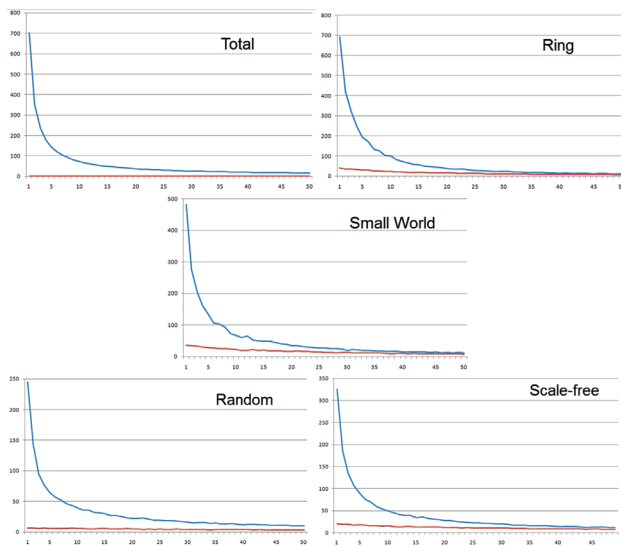


Fig. 6 Times to belief convergence in various networks for increasing links between sub-networks (shown in blue) and within single networks of that type (shown in red).

The horizontal axis in each case indicates added links between 1 and 50. The vertical axis varies in scale. In the case of total and ring networks it extends from 0 to 800 steps; in the case of small worlds it extends to 500 steps, with only 250 and 350 in the case of random and scale-free networks. The two facts that we want to emphasize, however, are (a) the extreme divergence between rates of belief convergence between linked sub-networks and single networks in each case, and (b) the remarkable similarity of the curves for linked sub-networks in each case. We emphasize that similarity by plotting results for all sub-network types in log-log form in Figure 7.

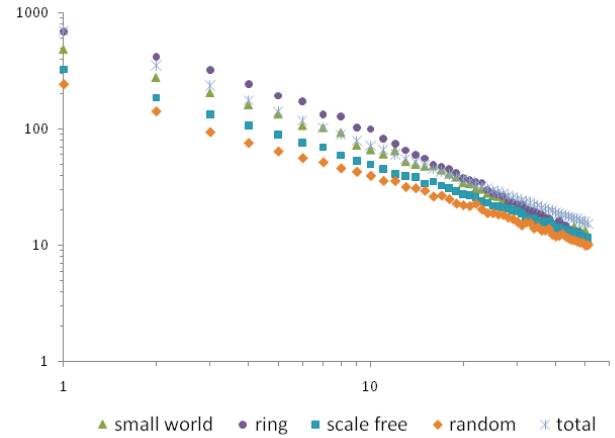


Fig. 7 Log-log plots of times to belief consensus with increased linkages between sub-networks of various types

Where information transmission is a matter of memes rather than germs, linkage degree effects follow the same pattern regardless of the structure of sub-networks. For meme transmission, unlike infection, it is degree of linkage that trumps network type. If one wants to plot the course of an epidemic, we noted, it is crucial that one knows the structure of networks involved. If one wants to plot the course of belief transmission, it is degree of linkages between sub-networks, of whatever type, that will be crucial.

As the differences in scale indicate, the particular structure of networks *is* important in order to gauge whether a single link between sub-networks will allow consensus in 250 steps or 700, as indicated for random and ring networks. The pattern of changes in belief transmission with increasing linkages between sub-networks from any initial point, however, is precisely the same regardless of network structure; the classic signature of power law distributions (Newman 2001, 2005).

The Dynamics of Genetic Transfer

Genetic information transfer is characterized by crossover from sexual reproduction, gene flow from different

populations, and random mutation of genetic material. To simulate this type of information transfer, we gave each agent 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 sub-networks, each sub-network begins with a uniform genetic code of this type. In the case of single networks, we randomize the two codes across agents.

On each time step of the model, each agent pairs off with an unpaired agent she is connected to, if there is such an agent. Each pair then mates. Two new genetic codes are formed, each of which consists of code from one parent to the left of a random crossover point and code from the other parent to the right. The production of the two codes only differs in the location of the random crossover point. Parents are replaced by offspring with the new genetic codes but with the same network connections. The process can be thought of as sexual reproduction within a population in which agents are replaced each round by their successors.

In the limit, this form of genetic updating will give us convergence of genetic codes in connected portions of networks. As the network converges, genetic information becomes more uniformly spread across the community, and we can therefore use time to convergence as a measure of the dynamics of genetic information.

An easy way to measure convergence in strings of 100 binaries would be to sum the ones of all strings and count the set of strings as converged when the sum for each string is within some chosen bounds. Though computationally cheap, however, such a measure would not adequately capture the idea we're after. On this measure, an agent with a genetic code of fifty ones followed by fifty zeroes would be maximally close to one with fifty zeroes followed by fifty ones, though in fact those codes differ at every point. A much better measure of genetic difference turns out to be Hamming distance--the number of places in which two codes differ.

Total convergence of a network might be thought of as that point when total Hamming distance is zero, i.e. when all agents have the same code. Given the stochastic nature of our simulation, we use instead time to a certain Hamming threshold. In particular, we take as our convergence measure that point at which two agents drawn randomly from the population can be expected to differ in less than 20% of their genetic code.⁸

If we think of germs, genes, and memes as strategies of information transfer across networks, which is the transfer of genetic information more like? Do the dynamics of genetic information transfer more closely resemble those of infection, or of belief?

⁸ In its current form our genetic model does not implement mutation, which we expect would merely increase the stochasticity of results. Issues of selection are among those raised in the conclusion as important areas for further work.

The answer is that genetic mixing doesn't match either of the other patterns in all respects, though it exhibits intriguing features of each.




Fig. 8 Generations to genetic convergence with a Hamming distance of 20% for increased linkages between sub-networks of various types.

Figure 8 shows genetic dynamics results for linked sub-networks of each of our types. Here, as in the case of belief, it is clear that network type makes very little difference. For reasons we don't yet understand, our scale-free preferential attachment networks are the outlier. When that case is removed, the proximity of results for increased linkages and regardless of the types of sub-networks linked is even clearer. Those results are shown in a re-scaled graph in Figure 9 and in log-log form in Figure 10.




Fig. 9 Generations to genetic convergence with a Hamming distance of 20% for increased linkages between total, ring, small world and random sub-networks.




Fig. 10 Log-log plot of genetic convergence for increased links between total, ring, small world and random sub-networks.

Here, as in the case of belief, we have the signature of a power law, though the slope or scaling exponent is very different. For genetics as for belief change--for genes, like memes--the differences in types of linked networks are of relatively minor importance.

In the case of belief, however, there was a remarkable difference between results for linked sub-networks and for single networks of a given type. In the case of infection there was not; dynamics for infection on linked networks of a given type very much paralleled those for single networks with the same number of added linkages.

This aspect of infection dynamics reappears for the transfer of genetic information as well. Figure 11 shows comparisons for our graph types between added linkages between networks (shown in blue) and within a single network of the same type (shown in red).

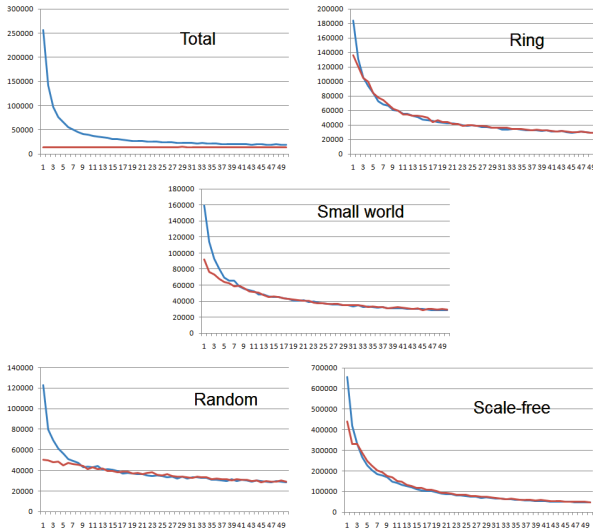


Fig. 11 Times to genetic convergence in various networks for increasing links between sub-networks (shown in blue) and within single networks of that type (shown in red).

In Figure 11 only the case of total networks mirrors the radical difference in single and linked networks characteristic of belief change, and that is because 'added' links in a total network are simply redundant. In all other cases single networks start with fewer generations to convergence, but after a very few number of added links the number of generations for single and linked sub-networks track nearly identical paths. In this regard the spread of genetic information is more like that of infection than belief change.

So do genes operate across networks like germs, or like memes? The answer is that genes operate entirely like neither germs nor memes, but they do operate a bit like each. In the case of belief transference it is linkage between sub-networks that is crucial, with types of networks linked of lesser importance. In measuring time to total infection, it is network type that is far more important than degree of linkage. In the transfer of genetic information, both aspects of networks--type and linkage--appear to be swamped in importance by the operation of the mechanism itself. Increased linkages certainly do increase genetic spread, but in ways that do not depend on either the type of network at issue or whether those linkages are within a single network or between sub-networks.

It is important to emphasize that these are early results, suggestive but in need of further exploration.⁹ What the results suggest, however, is that the transfer dynamics of genetic information may be surprisingly robust across network differences, insensitive to network factors that play a major role in the transfer of infection and belief.

Conclusions and Future Work

Our focus here has been on contrasting dynamics of transmission across networks, with an eye in particular to the case of transmission across links between sub-networks.

For both people and animals, communities of interaction are composed of sub-communities with varying degrees of contact and isolation. Different types of network are at issue for different forms of transmission, of course; it is contact networks that are important for disease transmission, sexual or genetic networks that are at issue in genetic recombination, and communication networks that facilitate belief transmission and change in people. What we have tried to trace are the dynamic signatures of different kind of transfer on a range of networks.

⁹ It will be important to see how these results on genetic transfer hold up for modeling variations that incorporate mutation, different forms of encoding of different lengths, and different patterns of generation replacement, for example.

In the abstract, each of these might be considered a form of *information* transfer. Belief change is a clear paradigm of information transfer, but it is also common to speak of transfer of genetic information. We can thus speak of the transfer of genetic information across networks of sexual mixing or hybridization. The spread of a disease across a network can be seen as the spread of information contained in the pathogen at issue.

Thought of in these terms, what our results indicate is that information dynamics are not all alike. The specifics of information transfer on networks will depend crucially on the type of information at issue and the specific mechanisms of transfer. Indeed what aspects of network structure are crucial for a particular kind of information transfer will depend on the particular kind of information and specific transfer mechanisms at issue.

If we think in terms of information, however, there is one major factor that is not included in our models here. Tracking that further aspect is planned as the next step in our research. Information produces certain effects and serves certain ends. Information in general, and each of the types of information considered here, is subject to selective pressure, typically conceived of in terms of fitness functions and payoffs.

A successful pathogen is one that effectively infects a host population so as to produce optimal reproduction and spread. But germs must stop short of annihilating their host population, which would spell their own doom as well. Successful interventions, on the other hand, are those that minimize or extinguish such a pattern of infection. What forms of intervention best target what forms of infection on what kinds of networks? What infection strategies, resistant to what forms of intervention, prove optimal for various types of network for each type of information under consideration?

Fitness functions are familiar in theoretical ecology and genetics. Our aim will be to incorporate selective pressure in terms of fitness in these models as well. How do selective pressures affect the different strategies of genetic recombination and asexual reproduction, for example, and on what kinds of networks? Are there particular environments, or environmental interfaces, that are most effectively exploited by linked sub-communities or gene pools? With an eye to both efficiency and resiliency, what genetic strategies prove optimal for what kinds of networks?

Just as genetics is tied to considerations of reproductive fitness, belief is tied to the selective pressures of truth and action. Here, in ways explored in some of our other work already, we want to investigate the dynamics of belief-formation strategies and the social networks on which they operate. Some of the questions we are after are descriptive, again linked to questions of public health: What belief interventions might be most effective in influencing beliefs and behaviors across segregated sub-communities with different levels of trust in particular

information sources?¹⁰ Other questions are normative: What network configuration of what kinds of epistemic agents, for example, with what kinds of belief-revision strategies, can best navigate a particular epistemic landscape?¹¹

Germs, genes, and memes can be considered families of information strategies played out on various forms of network. The next step will be to investigate which forms of such strategies optimize which forms of fitness on what kinds of networks.

Acknowledgements

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. Grim is grateful for his time with the Center for Complex Systems and Department of Philosophy at the University of Michigan.

References

- Auld, M.C., 2003. Choices, beliefs, and infectious disease dynamics. *Journal of Health Economics* 22, 361-377.
 - Barabási, A.-L. & Albert, R., 1999. Emergence of scaling in random networks. *Science* 286, 509-512.
 - Barrett, C. L., K. Bisset, J. Leidig, A. Marathe, & M. Marathe, 2009. Estimating the impact of public and private strategies for controlling an epidemic: A multi-agent approach. *Proceedings of the Twenty-First Innovative Applications of Artificial Intelligence Conference*, AAAI.
 - Börner, K., J. T. Maru, & R. L. Goldstone, 2003. The simultaneous evolution of author and paper networks. *Proceedings of the National Academy of Sciences of the USA*, 101 (Suppl.1), 5266-5273.
 - Centola, D., and M. Macy, 2007. Complex contagion and the weakness of long ties. *American Journal of Sociology* 113, 702-34.
 - Del Valle, S., H. Hethcote, J. M. Hyman, & C. Castillo-Chavez, 2005. Effects of behavioral changes in a smallpox attack model. *Mathematical Biosciences*, 195, 228-251.
 - Epstein, J. M., J. Parker, D. Cummings, and R. A. Hammond, 2008. Coupled contagion dynamics of fear and disease: Mathematical and computational explorations. *PLoS ONE* 3(12): e3955. doi:10.1371/journal.pone.0003955
 - Freeman, L. C., 1978. Segregation in social networks. *Sociological Methods and Research* 6: 411-429.
-
- ¹⁰ See Grim, Thomas, Fisher, Reade, Singer, Garza, Fryer, & Chatman 2011.
- ¹¹ See Grim, Reade, Singer, Fisher, & Majewicz, 2010a, 2010b.

Funk, S., E. Gilad, C. Watkins, & V. A. A. Jansen, 2009. The spread of awareness and its impact on epidemic outbreaks. *Proceedings of the National Academy of Sciences* 106 (16), 6872-6877, www.pnas.org/cgi/doi/10.1073/pnas.

Garber, T., J. Goldenberg, B. Libai & E. Muller, 2004. From density to density: Using spatial analysis for early prediction of new product success. *Marketing Science* 23, 419-429.

Goldenberg, J., B. Libai & E. Muller, 2001. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters* 12, 211-223.

Golub, B., & M. O. Jackson, forthcoming. How homophily affects learning and diffusion in networks. Available as a working paper at www.bepress.com/feem/paper296

Grim, P., C. Reade, D. J. Singer, S. Fisher, & S. Majewicz, 2010a. "Robustness across the Structure of Sub-Networks: The Contrast between Infection and Information Dynamics," *Complex Adaptive Systems: Resilience, Robustness, and Evolvability*, AAAI Technical Report FS-10-03, 2010.

Grim, P., C. Reade, D. J. Singer, S. Fisher & S. Majewicz, 2010b. "What You Believe Travels Differently: Information and Infection Dynamics Across Sub-Networks," *Connections* 30, 50-63.

Grim, P., S. B. Thomas, S. Fisher, C. Reade, D. J. Singer, M. A. Garza, C. S. Fryer, and J. Chatman (2011), "Belief Dynamics across Black and White Communities: A Data-Driven Agent-Based Network Model," submitted for publication.

Hallett, T. B., S. Gregson, J. J. C. Lewis, B. A. Lopman, & G. P. Garnett, 2007. Africa: Behavioral change in generalized HIV epidemics: Impact of reducing cross-generational sex and delaying age at sexual debut. *Sexually Transmitted Infections* 83, p. i50-i54.

Harrison, J. A., P. D. Mullen, & L. W. Green (1992). A meta-analysis of studies of the health belief model. *Health Education Research* 7:107-116.

Janz, N. K., & M. H. Becker (1984). The health belief model: A decade later. *Health Education Quarterly* 11:1-

Mullen, P. D., J. Hersey, & D. C. Iverson (1987). Health behavior models compared. *Social Science and Medicine* 24:973-981.

Newman, M. E. J., 2001. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98 (2), 404-409, reprinted in Newman, Barbási, & Watts, *The Structure and Dynamics of Networks*, Princeton: Princeton University Press, 2006.

Newman, M. E. J., 2005. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* 46: 323-351.