

Robustness Across the Structure of Sub-Networks: The Contrast Between Infection and Information Dynamics

Patrick Grim,^a Christopher Reade,^b Daniel J. Singer,^c Steven Fisher,^d
& Stephen Majewicz^e

^{abc}Department of Philosophy, Stony Brook University, Stony Brook, NY 11794

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

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

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

^eDepartment of Mathematics, Kingsborough Community College, Brooklyn, NY 11235

pgrim@notes.cc.sunysb.edu

Abstract

In this paper we make a simple theoretical point using a practical issue as an example. The simple theoretical point is that robustness is not 'all or nothing': in asking whether a system is robust one has to ask 'robust with respect to what property?' and 'robust over what set of changes in the system?'

The practical issue used to illustrate the point is an examination of degrees of linkage between sub-networks and a pointed contrast in robustness and fragility between the dynamics of (1) contact infection and (2) information transfer or belief change. Time to infection across linked sub-networks, it turns out, is fairly robust with regard to the degree of linkage between them. Time to infection is fragile and sensitive, however, with regard to the type of sub-network involved: total, ring, small world, random, or scale-free. Aspects of robustness and fragility are reversed where it is belief updating with reinforcement rather than infection that is at issue. In information dynamics, the pattern of time to consensus is robust across changes in network type but remarkably fragile with respect to degree of linkage between sub-networks.

These results have important implications for public health interventions in realistic social networks, particularly with an eye to ethnic and socio-economic sub-communities, and in social networks with sub-communities changing in structure or linkage.

Introduction

To ask whether a system is robust or resilient is to ask an incomplete question. Explicitly or implicitly, we are always interested in whether some specific properties of a system remain under some specifically envisaged sets of pressures or changes. In asking whether a system is robust or resilient, we are always asking 'robust with respect to

what?' and 'robust over what envisaged pressures or changes?'

Here we offer a particular case in which answers to questions of robustness across the structure of sub-networks are importantly different for different transfer dynamics. In a thumbnail sketch, subject to later qualifications, the contrast is this:

	Degree of Linkage	Sub-Network Type
Infection Dynamics	Robust	Fragile
Information Dynamics	Fragile	Robust

The dynamics of infection across linked sub-networks is relatively robust with regard to the degree of linkage. But it is importantly fragile with regard to the structure of the sub-networks themselves: total, ring, small world, random, or scale-free.

The dynamics of information, in contrast—where what is at issue is transfer of beliefs rather than diseases—reverses these characteristics. Given reinforcement effects, the dynamics of information is importantly fragile with regard to degree of linkage. The pattern of transfer, however, is notably robust across the structure of sub-networks, taking the characteristic form of a power law.

The focus of our particular example is public health issues regarding the ubiquitous phenomenon of social networks composed of heterogeneous sub-networks. Realistic social networks do not form a uniform and homogenous web. Social communities are composed of sub-communities, with varying degrees of contact and isolation between them both in terms of the physical contact necessary for disease transmission and the

informational contact crucial to transmission of belief. Racial, ethnic, socio-economic, demographic, and geographical sub-communities offer clear examples. Racial and socio-economic sub-communities may be more or less isolated or integrated with other sub-communities, with varying strengths not only of physical contact but of information transfer, communication, and trust. In the case of a pandemic, degree of isolation or integration will be crucial in predicting the course of contact and therefore the dynamics of disease transmission. But in such a case degree of informational isolation or integration will also be crucial in tracking changes in health care beliefs and behaviors, with both immediate and long-range effects on the course of the disease.

Public health has been one of the primary targets for agent-based and network modeling. A significant amount of work has been done on the role of network structure in the spread of disease (Meyers, Pourbohloul, Newman, Skowronski & Brunham 2005; Keeling 2005; Ferrari, Bansal, Meyers & Bjørnstad 2006; Miller & Hyman 2007; Eubank, Guclu, Kumar, Marathe, Srinivasan, Toroczkai, & Wang 2004). Significantly less has been done on the role of network structure in the dissemination of health beliefs and attitudes (see however Centola and Macy 2007 and Golub & Jackson, forthcoming).¹

It is clear, however, that 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).²

The dynamics of belief turns out to be very different from the dynamics of contact infection. For infection, measured in terms of average time to total infection across a network, it is the structure of the network or its sub-networks that is of primary importance—whether the basic network or networks at issue form rings, total networks, hubs, wheels, small worlds, scale-free or random networks. With respect to the structure of sub-

networks, in other words, time to infection is fragile. In gauging time to infection, degree of linkage between sub-networks is of relatively minor importance; with respect to degree of linkage, infection is surprisingly robust. The explanation for these results lies in the character of infection transfer—the fact that regardless of the number of connections to a node, a single contact is enough for infection to spread.

For belief transmission, in contrast, measured in terms of average time to total consensus, it is network structure that is of minor significance. The pattern of belief transmission forms a robust power-law regardless of the structure of sub-networks involved. Unlike infection, it is with regard to degree of linkage between sub-networks that the dynamics of belief is fragile. Here again the explanation lies in the character of the transfer at issue. Where reinforcement effects are at issue, the proportion of 'like-minded' connections to a node can be of major importance, with the result that nodes on connecting links play a decisive dynamic role.

Our effort here is to use this contrast between the dynamics of belief and infection across networks to emphasize the importance of asking not merely whether a system is robust but specifying what aspects of fragility and robustness are at stake with respect to what ranges of envisaged changes or pressures. More complete details of both analytic results and simulational results are available in an on-line appendix at www.pgrim.org/robustness.

I. Infection Dynamics Across Linked Sub-Networks

A. A First Example of the Importance of Structure: Ring and Total Networks

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. We will use degree of linkage in a relative sense to refer to increased c or bridges of this sort. 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.

¹ Centola and May 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 outline analytic results on homophily in random networks, with a similar emphasis on the contrast between diffusion and belief averaging. Our work here, part analytic and part simulational, extends that work and shows that the central contrast holds across networks of various types.

² Though not fully explored here, work on networks, diffusion of innovation, and marketing may also be relevant (Valente 1995; Goldenberg, Libai & Muller 2001; Garber, Goldenberg, Libai & Muller 2004; Watts & Dodds 2007).




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

We focus on varying degrees of connection between sub-networks of varying structure. For simplicity we use just two sub-networks of equal size, concentrating on ring sub-networks, total or connected networks, small worlds, random and scale-free sub-networks. How does the degree of connection between two sub-networks affect the dynamics of diffusion or infection across the network as a whole? How do results on degree of connection between sub-networks of a specific structure compare with results on a single network of the same structure to which the same number of links are added?

Some results here are simple and analytic, but also indicate the variety that can be expected. Consider, at one extreme, a network composed of two totally connected sub-networks with a single link between them, as in the second network in Figure 1. How many steps will be required to total infection, starting from a single random infected node? Assuming a 100% infection rate, where n is the total number of nodes, the average number of steps to total infection is

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

where n is the total number of nodes. From any node other than those on the ends of our connecting link, there are three steps to total infection: (1) to all nodes of the immediate connected networks, (2) across the one connecting link, and (3) from there to all nodes of the opposite connected network. If the initially infected node is one of those on the ends of our connecting link, there are merely two steps to total infection, giving us the formula above.

Adding further links has no dramatic effect in such a case; steps to total infection is robust across increased linkage. 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, again assuming an infection rate of 100% rate and incorporating n nodes and

m discrete links between sub-networks (links sharing no nodes),³ 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 $m = .5n$, time to infection approaches a limit of 2.

For a single total network, like that on the left in Figure 1, additional linkages will clearly have no effect at all: infection will in all cases be in a single step, and additional links will simply be redundant.

Where sub-networks are total, variance in infection time is necessarily just between 2 and 3 steps. At the other extreme is the case of a network with rings as sub-components. 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 $\frac{n}{4} + \frac{1}{2}$.

Variance between maximum and minimum times to total infection is therefore 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 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.

The structure of sub-networks is crucial for other factors as well. We have noted that increasing links

³ 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.

between sub-networks has a minimal effect where those sub-networks are total. Where sub-networks are rings of 50 nodes, in contrast, the effect is dramatic. The top line in Figure 2 shows results from a computer-instantiated agent-based model in which we progressively increase the number of links between random nodes of those sub-networks from 1 to 50. 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. 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.

Similar simulational results for added links between total sub-networks, in contrast, show a relatively flat result with decline in average time to infection from only 2.98 to 2.35. Difference in network structure clearly makes a major difference in time to total infection. That difference is not due to degree of linkage between sub-networks, however. A graph of results in which links are added not between ring sub-networks but across a single ring shows a result almost identical to that in Figure 2.

The lesson from ring and total networks is that it is not the degree of linkage between sub-networks that affects time to total infection but overall network structure itself, whether characterizing a single network or linked sub-networks. Changes in infection rates with additional random links (1) across a single network and (2) between two smaller networks with the same structure show very much the same pattern. Degrees of linkage between sub-networks interact with the structure of those sub-networks in order to generate patterns of infection, but it is the structure of the networks rather than the degree of linkage that plays the primary role. Analytical and simulational results for hub and wheel networks, very much in line with conclusions above, are available in an online appendix (www.pgrim.org/robustness)

B. Infection Across Small World, Random, and Scale-Free Networks

For patterns of infection, the importance of general structure type over degree of linkage between sub-networks holds for small world, scale free, and random networks as well. Results for small world networks are shown in the second line from the top in Figure 2, here with roughly a 9% probability of rewiring for each node in an initial single ring. (Our probability is 'roughly' 9% because in each case we add minimal links so as to assure a connected network. Without that assurance, of course, infection is not guaranteed to percolate through the network as a whole.) Increasing linkages from 1 to 50 results in a decrease in steps to total infection from 22.5 steps to 7.45. Increasing links within a single small world follows virtually the same pattern, with a decrease from 19.8 to 7.2.

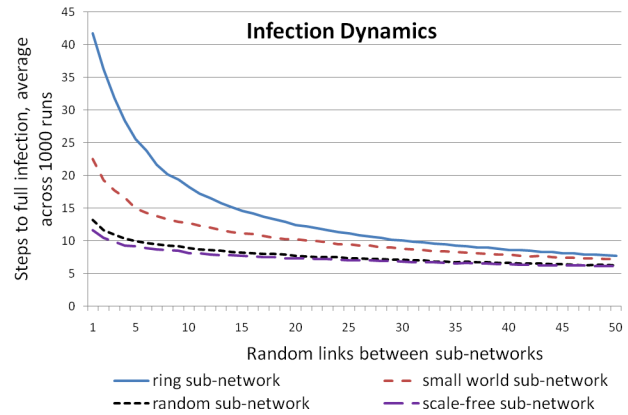


Fig. 2 Average time to total infection with increasing links between ring, small world, random, and scale-free sub-networks

Similar results for random and scale-free networks appear in the third and fourth graphed lines of Figure 4. For random networks, roughly 4.5 percent of possible connections are instantiated within each sub-network, with minimal links needed to guarantee connected networks. Our scale-free networks are constructed by the preferential attachment algorithm of Barabási and Albert (1999).

Here as before there is little difference where additional links are added within a single network, whether small-world or scale-free. In each case the number of initial steps is slightly smaller, but only in the first 10 steps or so is there any significant difference and convergence is to the same point. In the case of random networks, times decrease from 9.79 to 6.45. In the case of scale-free networks, times decrease from 7.9 to 6.08.

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. 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.

C. Infection on Networks: Qualifications and Provisos

Results to this point have been calculated with an assumption of 100% infection—a disease guaranteed to be transmitted at every time-point of contact between individuals. More realistic assumptions regarding rate of infection affect the rates calculated above, more pointedly emphasizing the importance of structure. Here we again use ring and total networks as an example.

Where sub-networks are total, probability of infection from single contact really makes a difference only at the link between sub-networks: as long as the probability of infection exceeds $2/n$, a quick infection of all individuals in the total sub-networks is virtually guaranteed. Simulation results indicate that with a single link between total sub-networks average time to full infection shifts only from an average of 3.8 steps to an average of 2.98 with a change of infection rate from 100% to 50%. For ring sub-networks, on the other hand, the same change in infection rate roughly doubles times to full infection across all numbers of linkages.

For more realistic infection rates, therefore, it is more important rather than less to know the structure of social networks. If those sub-networks approximate total networks, neither infection rate nor additional links between sub-networks make much difference. If sub-networks approximate ring networks, both number of links and infection rate will make a dramatic difference in the course of an 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 here we need to add an important proviso: 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 3 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.

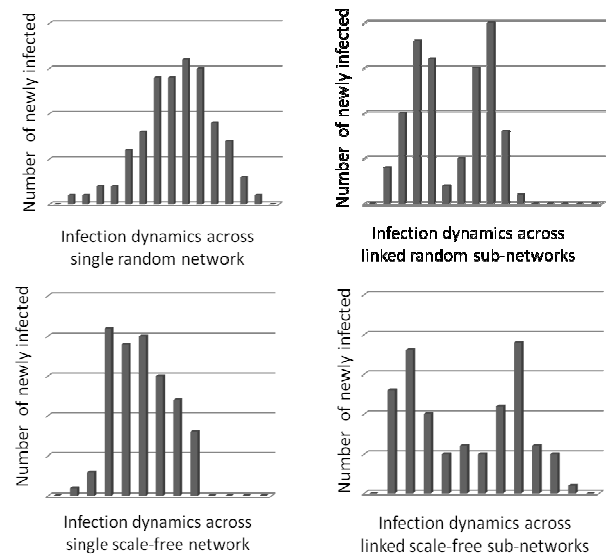


Fig. 3 Contrast in typical dynamics of infection in single and linked sub-networks, even where time to total infection is the same

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.

I. Information Dynamics Across Linked Sub-Networks

What you believe travels differently. In what follows we use a simple model of belief updating to show the crucial importance of degree of sub-network linkage in belief or information transmission across a network. Some earlier results have noted 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.

That beliefs will change in accord with some pattern of reinforcement along these lines is very plausible, backed by a range of social psychological data, and is therefore an aspect of realism in the model. What is unrealistic is the particular form of reinforcement instantiated here—the particularly simple pattern of belief averaging, applied homogeneously across all agents. In order to be informative regarding an exterior reality a model, like any theory, must capture relevant aspects of that reality. In order to offer both tractability and understanding a model, like any theory, must simplify. This first model of belief transmission is intended to capture a reality of belief reinforcement; the admittedly artificial assumption of belief averaging is our simplification.

Our attempt, then, 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. In particular, we want to emphasize the major differences between the dynamics of belief change across information networks and the dynamics of infection diffusion across contact networks, outlined above. What you believe travels differently.

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, 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.

A. Belief Diffusion Across Ring and Total Networks

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. As we add additional links between sub-networks, how does the dynamics of belief diffusion change, measured in terms of time to consensus across the community?

We progressively add random links (1) between belief-polarized ring sub-networks, and (2) within a single ring network of belief-polarized agents. Average times to consensus are shown in Figure 4.

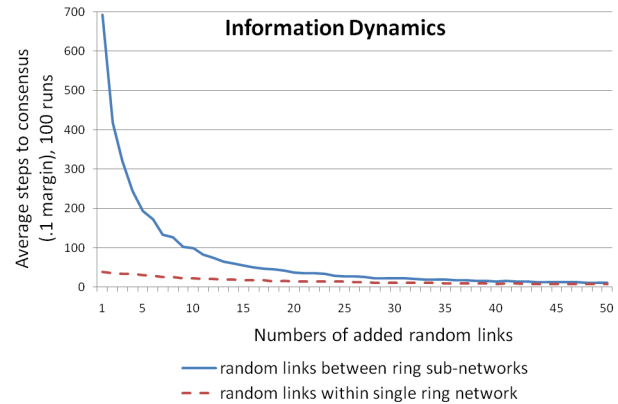


Fig. 4. The importance of degree of linkage in time to belief consensus: Contrast between results of adding additional linkages within a single ring network (below) and between ring sub-networks (above)

Increasing linkages between polarized ring sub-networks makes a dramatic difference. Time to consensus, unlike time to infection, is fragile with regard to degree of linkage. Average time to consensus for a single linkage in such a case is 692.44; average time to consensus for 50 linkages is 11.59, with a distinct and characteristic curve between them. For infection, we noted, there is virtually no difference between added links within a single ring network and added links between ring sub-networks. In the case of belief, in contrast, there is a dramatic difference between the two graphs.

Within a single total network, all agents will achieve a mean belief in a single step; additional linkages in such a case are merely redundant. Results for linked total sub-networks, in contrast, parallel those for rings above. Average steps to belief convergence with a single link approximate 700 steps in both cases; with 50 links, average time to convergence is 12 in the case of rings and 16 in the case of total sub-networks. The overall pattern of the two graphs is also very much the same. What that similarity shows is the strikingly robust effect of degree linkage in each case: an effect that in the transmission of belief overrides the fact that we are dealing with totally distributed ring networks in one case, totally connected networks in the other.

B. Belief Transmission across Small World, Random, and Scale-Free Networks

The same contrasts between single and linked sub-networks in the case of belief transmission hold for other network structures as well.

The effect of added linkages within a single small-world network closely parallels that for the single ring shown above. The effect of adding linkages between small-world sub-networks is again very different. In absolute terms results for small worlds differ from those shown for rings, declining from 481 steps to 11.4. The shape of the curve for small worlds, however, is very much that shown for rings above.

Given a single random network, using 2.25% of possible linkages, additional linkages give a decline in time to belief consensus from only approximately 6 steps to 4. Where random sub-networks are at issue (using 4.5% of possible linkages in each sub-network), the curve is again that displayed for rings above, though here absolute values decline from 244 to 10.15.

For single scale-free networks, additional linkages give a roughly linear decline from 20 to 7 steps. For scale-free sub-networks, additional linkages again follow the curve shown above, here with absolute values dipping from 325 to 11.73.

A similar curve characterizes effects of degree linkage in belief transmission regardless of the basic structure of the sub-networks involved: it is with regard to the shape of the curve that results are robust across linkage differences, despite differences in initial absolute values. We emphasize the robustness of belief transmission patterns by plotting results for all sub-network types in log-log form in Figure 5.

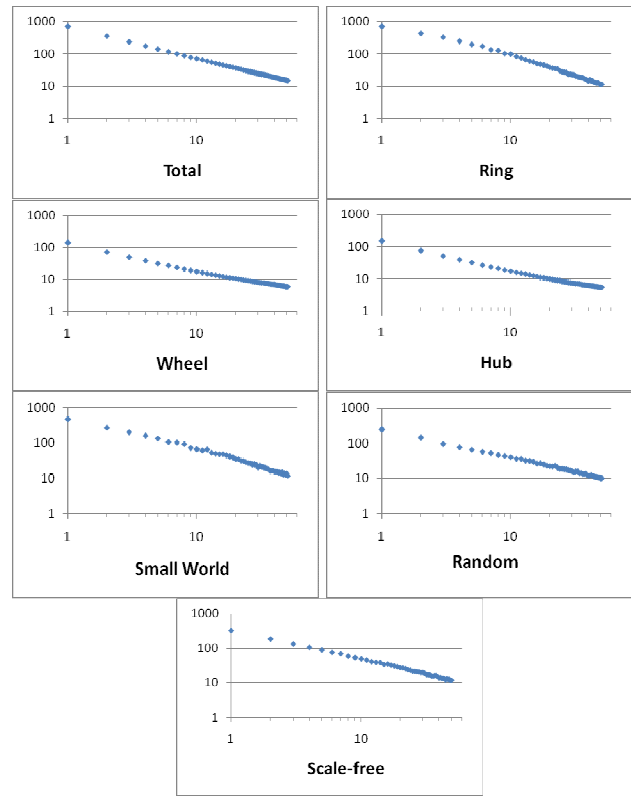


Fig. 5 Time to belief consensus with increasing linkages between sub-networks of various types, plotted log-log.

Linkage degree effects follow the same pattern regardless of the structure of sub-networks. If one wants to plot the course of an epidemic, we noted in section I, it is crucial that one knows the structure of networks involved. If one wants to plot the course of belief transmission, knowledge of structure is much less important.

The particular structure of networks is important in order to gauge whether a single link between sub-networks will allow consensus in 140 steps or 700, as indicated for hub and total networks above. 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. That pattern is the classic signature of power law distributions, indicating that the relationship between increased linkage and time to consensus parallels a range of natural and social phenomena, including the relationship between frequency and size of earthquakes, metabolic rate and body mass of a species, and size of a city and the number of patents it produces. Power law distributions also appear in some empirically observed characteristics of biochemical, protein, citation and sexual contact networks (Faloutsos, Faloutsos, & Faloutsos, 1999; Jeong, Tombor, Albert, Ottvai, & Barbási 2000; Fell & Wagner 2000; Liljeros, Edling, Amaral, Stanley, & Åberg 2001; Newman 2001, 2005). The fact that such an effect appears in linkage

effects on the dynamics of belief suggests the possibility of incorporating a range of theoretical and methodological work from other disciplines in studying behavior dynamics in the spread of disease, particularly with an eye to the effect of belief polarization, health care disparities, and social linkage or integration between ethnic and socio-economic sub-communities.

III. Conclusions and Future Work

Our focus here has been on the structure of contact and informational networks and the very different impact of aspects of that structure on the dynamics of infection and information. With respect to linkage between sub-networks, time to total infection is relatively robust whereas time to belief consensus is remarkably fragile. With respect to the structure within sub-networks, time to total infection is fragile whereas time to total consensus takes the robust form of a power law.

For infection, measured in terms of average time to total infection across a network, it is the structure of the network or sub-networks that trumps other effects. In attempting to gauge time to total infection across a community, therefore, the primary piece of information needed is whether the social network or component networks at issue approximate rings, hubs, wheels, small worlds, random, scale-free or total networks. For time to total infection, degree of linkage between sub-networks is of much less importance, though we have noted that points of linkage continue to play an important role with regard to fragility and prospects for targeted intervention.

For information, measured in terms of average time to belief consensus, the importance of general structure and linkage between sub-networks are reversed. On the model of belief used here, in attempting to gauge the dynamics of information flow across a community, the primary piece of information needed is the degree of linkage between composite sub-communities, whatever their internal structure. The fact that the particular structure of those sub-communities is of lesser importance is highlighted by the fact that average time to belief consensus given increasing linkages follows the same familiar power-law pattern regardless of networks structures involved.

The explanation for each result, and for their contrast, lies ultimately in the character of the transfer itself. Here we can only sketch such an explanation, though the essentials are fairly intuitive. Transmission of infection demands only a single connection from node to node. If sub-network structures vary importantly in structure and density of nodal connections, the result will be a major difference in disease transmission, often swamping transmission delay across linkages between sub-networks.

Because of reinforcement effects, transmission of belief often demands more than a single connection. A node situated at a linkage between polarized sub-networks will

typically have many connections on one side, repeatedly outweighing input from a single node on the other side. Where reinforcement effects are at issue, therefore, linkages between sub-networks will play a major role, swamping the importance of network structure within the linked sub-networks.

It is quite plausible that belief transmission involves strong reinforcement effects; the model of belief used here is designed to capture such an effect. In other regards, however, the belief model is quite clearly artificial. In our work to this point, belief change is by simple averaging of information contacts, and all agents follow the same formula for belief updating. Our attempt in future work will be to test the robustness of conclusions here by considering a range of variations on that central model of belief change.

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, and by a pilot grant for Developing an Agent-Based Model to Assess Racial Differences in Medical Discrimination, Social Support and Trust, 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.
- Barbasi, 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
- Eubank, S., H. Guclu, V. A. Kumar, M. Marathe, A. Srinivasan, Z. Toroczkai, & N. Wang, 2004. Modeling disease outbreaks in realistic urban social networks. *Nature* 429:180-184.
- Fell, D., & A. Wagner, 2000. The small world of metabolism. *Nature Biotechnology* 18, 1121-1122, reprinted in Newman, Barbási, & Watts, *The Structure and Dynamics of Networks*, Princeton: Princeton University Press, 2006.
- Faloutsos, M, P. Faloutsos, & C. Faloutsos, 1999. On power-law relationships of the internet topology. SIGCOMM '99, reprinted in Newman, Barbási, & Watts, *The Structure and Dynamics of Networks*, Princeton: Princeton University Press, 2006.
- Ferrari, M. J., S. Bansal, L. A. Meyers & O. N. Bjørnstad, 2006. Network frailty and the geometry of herd immunity. *Proc. R. Soc. B* 273, 2743-2748
- Freeman, L. C., 1978. Segregation in social networks. *Sociological Methods and Research* 6: 411-429.
- 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.
- 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.
- Jeong, H., B. Tombor, R. Albert, Z. N. Ottvai, & A.-L. Barbási, 2000. The large-scale organization of metabolic networks. *Nature* 407, 651-654, reprinted in Newman, Barbási, & Watts, *The Structure and Dynamics of Networks*, Princeton: Princeton University Press, 2006.
- Keeling, M., 2005. The implications of network structure for epidemic dynamics. *Theoretical Population Biology* 67, 1-8.
- Liljeros, F., C. R. Edling, L. A. Nunes Amaral, H. E. Stanley, & Y. Åberg, 2001. The web of human sexual contacts. *Nature* 411, 907-908, reprinted in Newman, Barbási, & Watts, *The Structure and Dynamics of Networks*, Princeton: Princeton University Press, 2006.
- Meyers, A. M., B. Pourbohloul, M. E. J. Newman, D. M. Skowronski & R. C. Brunham, 2005. Network theory and SARS: predicting outbreak diversity. *Journal of Theoretical Biology* 232, 71-81.
- Miller, J. C. & J. M. Hyman, 2007. Effective vaccination strategies for realistic social networks, *Physica A* 386, 780-785.
- 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.
- Redner, S., 1998. How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B* 4, 131-134.
- Trotter, R. T., R. B. Rothenberg, & S. Coyle, 1995. Drug abuse and HIV prevention research: Expanding paradigms and network contributions to risk reduction. *Connections* 18, 29-45.
- Valente, T. W., 1995. *Network Models of the Diffusion of Innovations*. Creskill, N.J.: Hampton Press
- Watts, D. J., & P. S. Dodds, 2007. Influentials, networks, and public opinion formation. *Journal of Consumer Research* 34 (2007), 441-458.