Concurrency and Commitment:
Network Scheduling and its Consequences for Diffusion*

David R. Gibson

Harvard University

536 William James Hall

33 Kirkland St.

Cambridge, MA 02138

dgibson@wjh.harvard.edu

* For comments on an earlier version of this paper I am indebted to Jim Ennis and Duncan Watts. For their suggestions on the model, I am grateful to Freda Lynn and participants in my graduate Modeling Social Dynamics class.

Abstract

People are embedded in networks involving many relationships that exist concurrently. They are not, however, enacted concurrently: because a person can only be in one place at a time, and by and large one tie can only be enacted insofar as others are not, the interactional obligations that concurrently adhere to a network position must be taken up sequentially. Complicating matters further, face-to-face encounters require that two people be *simultaneously* available. From these circumstances arise the conventions, strategies, and experiences of scheduling. Here I use a computer simulation to model how networks are concretely enacted through dyadic encounters, and explore the implications of this for diffusion. Whether, and how quickly, a practice (or idea, or disease) spreads through a network is shown to depend upon the interaction of network structure, whether or not actors try to compensate for past scheduling imperfections, and whether actors adopt upon first exposure, or upon finding that some proportion of their recently encountered friends have already done so. The study of network scheduling is one component of a more general program for the study of "enactment dynamics," or the way in which pre-existing ties get activated and deactivated under concrete circumstances, with important consequences for what are counted as "network effects."

Introduction

A person's network is lived, and experienced, as an irregular succession of planned engagements, chance meetings, extended lulls, unanswered emails, I-was-just-about-to-call-yous, postponements, reschedulings, and competing priorities. It is represented and analyzed, however, very differently: as a static configuration of present-or-absent ties that exist concurrently, even when that means that a given person is portrayed as equally involved in several relationships simultaneously. Odd as this seems, or should seem, as an opening methodological gambit, it yields immediate advantages, clearing the way for the rigorous operationalization of structural properties such as density, centrality, and path distance (Wasserman and Faust 1994), and for the application of a slew of graph-theoretic techniques (Albert and Barabási 2002; Pattison and Wasserman 1999; Strogatz 2001).

More significantly, networks thus construed have been linked to a range of non-network outcomes, such as innovation diffusion (Burt 1987; Galaskiewicz and Burt 1991), managerial promotion (Burt 1992; Podolny and Baron 1997), organizational membership (McPherson, Popielarz, and Drobnic 1992), and health (e.g., Cohen, Doyle, Skoner, Rabin, and Gwaltney 1997). It appears, then, that network representations capture something important about the world. As a first approximation, it may be that we are continuously pulled into the orbit of our relationships, so that given enough time, a person encounters everyone in his or her network, with the result that he or she makes decisions, or suffers the decisions of others, *as if* on the basis of relationships simultaneously in play, as network imagery implies.

This, however, is difficult to reconcile with some rather obvious facts about the world. In particular, scheduling considerations – related to who is available to interact with whom at any point in time – may prevent a person from consulting with all of his or her advisors prior to an important decision, or from gathering information from all of his or her acquaintances before changing jobs, or from garnering significant emotional support during a period of personal crisis. Thus the premise of this paper: *scheduling rules and constraints interpose between social networks and the outcomes we attribute to them*, including those listed above. The results of ignoring this fact are two-fold: first, the mechanisms behind known network effects are likely to be mis-specified (e.g., as being mainly psychological), and second, some network effects are likely to be entirely missed – specifically, when the play of scheduling across a network produces results unrecognizable from the perspective of the network taken alone.

Thus it is important that we begin to understand how scheduling rules interact with networks to produce outcomes. Here I take one particular outcome of long-standing interest to network analysts, diffusion, and consider how the pace and extent of diffusion (of an idea, disease, or practice) is affected by the interaction of scheduling rules and network structure. After describing past work on network scheduling, I propose a computer simulation of the same, one which aims at some degree of phenomenological realism. The initial diffusion analysis assumes that contagion occurs the instant someone who has not adopted (or is not infected) encounters someone who has (or is). Then I explore contagion that occurs at a "threshold," when some fraction of one's friends have adopted, though I propose a modification of standard threshold models to incorporate the fact that scheduling constraints may prevent a person from accurately assessing this fraction. For both types of contagion, I explore the effects of network structure, and whether or not people deliberately try to correct for past scheduling imperfections, which I call "remediation," on time to complete diffusion. In the Conclusion I argue that the study of network scheduling is one branch of a wider research program into "network enactment," which is concerned with the situational activation and deactivation of network ties, and the consequences of this for all of what get counted as "network effects."

Network scheduling

The mere fact that I can only be in one place at a time imposes severe constraints on my ability to enact my network ties through dyadic encounters – through having lunch with friends, coffee with colleagues, meetings with advisors and students, etc. – and the problem is compounded many-fold by the fact that not only must I find time for a friend, he or she must also find time for me, and, moreover, these "times" must coincide. Though fundamental, these facts have received little attention in the sociological literature, though at least two articles have come close. Winship (1978), for one, devised an ingenuous equilibrium model for dyadic negotiations over time allocations, that is, over how much time two people will spend with one another given the constraint that Al cannot spend more or less time with Betty than Betty spends with Al. A limitation of his model, however, is that it ignores the fact that for two people to spend time together, they have to agree on when, precisely, that is to happen. Consider the case of three people, Al, Betty, and Carl, each of whom wants to spend one half of his or her time with each of the others. Because there is perfect symmetry in their preferences, this is exactly what Winship's model predicts will happen. The problem is that if Al spends the first half of each day with Betty, and the second half with Carl, there will be no point at which Betty and Carl will be simultaneously free to see each other.[2] This difficulty arises because Winship does not require his actors to commit to particular time slots, only to quantitative allocations.

If the weakness of Winship's approach is that time is detached from schedules, its strength is that negotiations are conducted in a decentralized manner, without top-down oversight – which is how most such decisions are actually made. Leifer (1990) takes the opposite approach, imposing real-time scheduling constraints as well as an omniscient central scheduling authority to contend with them. Because his main application is to the scheduling of games in professional sports, however, there is no pre-existing network awaiting enactment – the scheduling objective is to ensure a balance of home and away games. Further, the scheduling phenomenology is far from realistic for the types of scheduling scenarios that concern me here, and which concern Winship. For one thing, if the central authority in Leifer's model determines that a partially-completed schedule is destined to fall short of the scheduling objective, it can erase the encounters (games) already planned and start from scratch. Real people, however, do not have the option to pull out of their obligations whenever it suits them, and generally abide by their commitments even when the resulting schedules appear ill-conceived.

Among the features of a realistic model of scheduling, then, are (1) that it is performed in a decentralized manner, *contra* Leifer but as per Winship; and (2) that actors are concerned with scheduling specific encounters, rather than general allocations, *contra* Winship but as per Leifer. It must also (3) take a pre-existing network as input, so that people give scheduling priority to those they have pre-existing ties to.

Scheduling simulation

Simulations are increasingly used within sociology (Macy and Willer 2001), and the social sciences generally (Axelrod 1997; Gilbert and Troitzsch 1999), to study the macro consequences of variations in the local rules and constraints governing the behavior of individual actors. My

---

[2] Winship (1992) acknowledges as much in a later paper, in which he uses graph coloring theory to analyze the scheduling of group gatherings such as colloquia.

use of simulations differs from what is typical in two respects. First, I am more concerned than most simulation researchers with phenomenological realism, that is, with programming actors to abide by rules that empirical research or, at least, commonsense suggests are actually operating in the real world. Importantly, this does not, at least in the current context, come at the expense of model simplicity, which simulation proponents such as Macy and Willer (2001) rightly value, since the rules that people actually follow tend to be fairly simple solutions to social dilemmas that simulated actors also face – here, the problem of how to credibly commit to invitations and engagements given that everyone has an interest in keeping their options open for as long as possible.

Second, and relatedly, simulation researchers often downplay the extent to which their simple models speak to what happens in the real world, appealing instead to the value that simulations have as virtual experiments (Axelrod 1997), which at most tell us what sets of conditions *would be sufficient* to produce some outcome in the real world were they met (Epstein and Axtell 1996). My position, in contrast, is that simulations should be built upon things we know for certain about the world, such as that no one can be in two places simultaneously, so as to explore the consequences of these known facts. Of course, in the real world there may be other forces at work offsetting these consequences, but knowing what these consequences would be in a simple model then directs our attention to these other forces, which can then be incorporated into later model iterations. This approach is illustrated below, when I permit actors to compensate for past scheduling imperfections by deliberately pursuing alters they have infrequently encountered – this as one way in which people cope with the fundamental limitation that being in one place means not being somewhere else, that attending to one tie normally means temporarily neglecting others.

The model

The simulation "core," upon which the diffusion analysis will build, is diagrammed in Figure 1.[3] It takes as input a "preference" network. For the sake of convenience, I often equate the ties in this network with those of friends, or else speak of people "liking" one another, though the analysis extends to any network of background propensities to interact with some people rather than with others, such as in an organizational setting where this may be determined by task

[3] The program was written in the programming language C.

interdependence rather than by sentiment. I assume that the network is symmetrical, such that if *i* considers *j* a friend, *j* considers *i* a friend as well, and that ties are dichotomous, so that one person either considers a second a friend or does not. As with most simplifying assumptions in simulation research, these are made to facilitate analysis of the dynamics arising most directly from the mechanism of interest – here, scheduling. That being said, future research may experiment with valued and nonsymmetrical relations.

[FIGURE 1 ABOUT HERE]

I refer to the main scheduling unit with which people are assumed to be working as a "round." At the beginning of the round (1), each person sends out a single invitation to someone they like; in the baseline model, he or she selects this person at random from the set of his or her friends. These invitations are issued simultaneously. Any two individuals who have sent one another reciprocal invitations are then immediately paired off for an encounter (2), and each declines any remaining incoming invitations (3), which are immediately withdrawn. Then, the simulation returns to step (1), and anyone just turned down sends out a new invitation, avoiding the person (and later, the set of persons) whom he or she has already been turned down by; anyone still awaiting a reply takes no further action at this point. Reciprocal invitations again result in pairings and rejections (2, 3), and anyone who sent an invitation to some other previously-committed person is also turned down (3). Once again, the recently rejected send out a new set of invitations (1).

[FIGURE 2 ABOUT HERE]

The invitation-pairing-rejection cycle (1-3) repeats until no additional pairings are possible (4b), *or* until the entire system freezes (4a), with everyone still in the running waiting upon someone else, in what I refer to as a "contingency cycle." Figure 2 illustrates some ways in which this can happen. In each case, there are three or more people caught in a cycle of invitations, such that, for instance, Al is waiting on Betty, Betty is waiting on Carl, Carl is waiting on Diane, and Diane is waiting on Al. Further, others can become caught up in such cycles, if they are waiting on someone in the cycle, or on someone who is waiting upon someone in the cycle, and so forth. When this occurs, no further progress can be made. My solution is to randomly select one person from among those so affected to become “impatient,” with the additional requirement that she has at least one incoming offer. She then accepts one of these – in the baseline model, at random – and rescinds her outstanding invitation, following which the newly paired individuals decline other invitations as usual. The simulation then returns to step (1), continuing as before until another contingency cycle is encountered, or until no further encounters can be scheduled given the network and the specific scheduling rules at work (on which more shortly), at which point scheduling efforts for that round conclude. The scheduled encounters then occur, and the actors begin the whole process over again at the start of the next round.

While some features of this simulation were implemented for the sake of programming convenience – in particular, the simultaneous issuing of each new set of invitations – several others are intended to capture real features of scheduling phenomenology. Firstly, people discriminate between those with whom they want to spend time and those with whom they do not. Secondly, two people have to agree to meet before an encounter can occur – an obvious requirement that is, however, ignored in much simulation research (e.g., Skyrms and Pemantle 2000).[4] Thirdly, people are limited to one outstanding invitation at any time, with respect to any given encounter opportunity. While this is something of a simplification, it remains true that we rarely send out multiple invitations with respect to the same time slot, lest both parties prove available and we have to take back the invitation after it has been accepted. Fourthly, an invitation extended is rarely rescinded unless rejected. The reason, perhaps, is that invitations amount to a kind of “line,” in Goffman’s sense of "a pattern of verbal and nonverbal acts by which one expresses one's view of the situation and through this his evaluation of the participants, especially himself" (1967, p. 5). In this context, the “view” expressed pertains to the inviter’s assessment of the social desirability of the invitee. To rescind such a line is severely face-threatening, and risks the animosity of the person so slighted and the tarnishing the reputation of the person responsible. Fifthly, people feel bound to the appointments they have committed to. While this is also a simplification, it is, once again, largely true, and on those rare occasions on which we do cancel an appointment we are at special pains to give an account of ourselves.

---

[4] This is not to deny the fact of chance encounters, though I do not consider them here.

Lying between expedient and realistic is the assumption that encounters are dyadic, and the provision for impatience. Regarding the first, while encounters do sometimes involve more than two people, there are good reasons for people to prefer dyadic encounters. One is that these are more easily scheduled, since the greater the number of people to be accommodated, the more difficult the scheduling problem. Another is that dyadic encounters impose less of a cognitive/strategic load upon interactants, who only need to keep track of what one other person already knows, or can safely be told. Finally, even purely social encounters suffer when they involve more than two people. In the words of Emerson, “Two may talk and one may hear, but three cannot take part in a conversation of the most sincere and searching sort” (qtd. in Crystal and Crystal 2000, p. 140).

While impatience certainly does, in the real world, serve to re-start scheduling when it has stalled, its implementation here is based more on programming simplicity than on realism. In particular, at present someone becomes impatient only when the entire group has run aground scheduling-wise, which is not something that a real person can be expected to know. As with so many aspects of the simulation, future research may experiment with alternatives, such as endowing each person with a level of characterological impatience, and thereby the capacity to rescind an outgoing offer and accept an incoming offer after a certain period of waiting, regardless of what is happening to others. This would complicate the current model, however, and indeed create further complications still, which I take as best avoided for current purposes.

Remediation

The baseline model assumes that actors have weighted their network options equally, which is why, at steps (1) and (4a) in Figure 1, their selections – of a friend to whom to send an invitation, or of a friend from whom to accept an invitation when impatient – are made randomly. But it is possible that actors deliberately compensate for scheduling imperfections experienced in prior rounds by more vigorously pursuing those friends whom they have not recently encountered, or weighing more favorably invitations from such people. I refer to this as “scheduling remediation,” and it plays an important role in what follows. The calculations involved take as input the cumulative history of dyadic encounters as these have occurred over a number of rounds. This is compared to what would have been ideal from the perspective of each person.

Then individuals issue invitations, and consider invitations from others, on the basis of the difference between the ideal and the reality – basically, the ideal-reality "residual."

The "ideal" number of times a person *i* would like to have seen a person *j* as of round *r* is calculated as the normalized strength of the tie, ranging from 0 to 1 and summing to 1 across all of *i*'s alters, times *r*:

$$\hat{f}_{ijr} = r\left(\frac{x_{ij}}{x_{i+}}\right) \qquad \text{(Eq. 1)}$$

where $\hat{f}_{ijr}$ is how often *i* and *j* would ideally have met, from *i*'s perspective, after *r* rounds; x_{ij} = 1 if *i* likes *j*; and x_{i+} is *x*'s outdegree (summing across all alters, inclusive of *j*). Thus, if John likes four people, he is "expected," if he gets his way, to see each of them one-quarter of the time, or on twenty-five occasions as of round 100. If, on the other hand, he only likes one person, he is expected to impose himself on that person every round. Note that $\hat{f}_{ijr}$ is always defined from the perspective of *i*, which may or may not accord with some *j*'s view on the matter.

When scheduling remediation is operating, the probability of *i* sending an invitation to *j* is then calculated as:

$$\Pr(i\mathrm{V}j) = \begin{cases} \dfrac{x_{ij}}{x_{i+}} & \text{if} \quad r = 0 \\[2ex] \dfrac{\hat{f}_{ijr+1} - f_{ijr} + \left|\hat{f}_{ijr+1} - f_{ijr}\right|}{2\sum_{k=1}^{g}\left(\hat{f}_{ikr+1} - f_{ikr} + \left|\hat{f}_{ikr+1} - f_{ikr}\right|\right)} & \text{if} \quad r > 0 \end{cases} \qquad \text{(Eq. 2)}$$

where f_{ijr} is how many times *i* and *j* have met so far, after *r* rounds, $\hat{f}_{ijr+1}$ is how often i and j would ideally have met *after this round*, from i's perspective (and similarly for f_{ikr} and $\hat{f}_{ikr+1}$), and *g* is the number of people in the group. In the first round, when *w* = 0, *i* has the same probability of sending an invitation to each *j* for whom x_{ij} = 1. In subsequent rounds, however, *i* is more likely to send an invitation to someone he or she has seen too little of ($\hat{f}_{ijr}$ - f_{ijr} > 0), where this likelihood increases with the size of the residual, but has a *zero* probability of sending an invitation to someone he or she has seen enough of, or more than enough of ($\hat{f}_{ijr}$ - f_{ijr} ≤ 0), given the number of rounds that have passed so far (the absolute values in the numerator and denominator setting Pr(iVj) to 0 when this is the case). The ideal value is indexed to r+1, rather than to r, because otherwise there is a chance that someone will be put into a temporary stupor by the bottom half of Eq. 2, upon finding that they have encountered each person in their network the ideal number of times. Basically this keeps everyone slightly nervous, as a spur to sending out a new invitation even when a person has succeeded fabulously (and improbably) in enacting his or her network up to that point.

A similar equation is used to calculate i's probability of accepting an invitation from j given that i has an outstanding offer to l with regard to which i has become impatient:

$$\Pr(i \text{ accepts } j \mid i\text{V}l,\, j\text{V}i, \text{imp(i,l)}) = \begin{cases} \dfrac{x_{ij}}{x_{i+} - x_{il}} & \text{if} \quad r = 0 \\ \\ \dfrac{\hat{f}_{ijr+1} - f_{ijr} + \left| \hat{f}_{ijr+1} - f_{ijr} \right|}{2 \sum_{k=1, k \neq l}^{g} \left(\hat{f}_{ikr+1} - f_{ikr} + \left| \hat{f}_{ikr+1} - f_{ikr} \right| \right)} & \text{if} \quad r > 0 \end{cases} \qquad \text{(Eq. 3)}$$

where the notation is as in Equation 2, the difference being that the normalizing divisor here excludes l whom i has written off as a lost cause. If i has no invitation from a j whom i has seen too little of, i randomly selects an invitation to accept, as if r = 0.

Remediation changes the question people ask from, "Whom do I like?", to, "Of the people I like, whom haven't I seen much of?"[5] Note that I take remediation, or its absence, as a feature of a scheduling culture, so that the entire group is modeled as remediative or unremediative. While this elides the possibility that people differ *within* a group – that is, of heterogeneous scheduling strategies – it is a simplifying assumption that, arguably, maps, at least crudely, onto the world. It is easy to imagine, for instance, that business and professional networks are, on whole, more remediative than adolescent networks. Still, future research would do well to consider the possibility of heterogeneous strategies, both as the degree of

[5] As remediation is currently implemented, it turns people into perfect record-keepers, who care as much about what happened, or did not happen, a long time ago as what happened more recently. An alternative would be to add a tunable memory decay parameter, whereby people would assign more weight to recent encounters.

heterogeneity affects group-level outcomes such as diffusion, and as particular individuals are affected by their strategies and the strategies of their network alters, say in terms of their respective levels of social engagement.

Diffusion

Diffusion researchers normally take ties to operate concurrently. This is particularly true in research on the spread of an idea or practice.[6] In empirical research this means asking whether a person is more likely to acquire the practice if he or she has friends who already have (e.g., Burt 1987; Coleman, Katz, and Menzel 1967),[7] while computational research takes the answer to be yes, in order to explore the effects of various parameters – such as the degree of connectivity, and the proportion of one's friends that need to adopt before one decides to jump on the bandwagon – on the pace and ultimate extent of diffusion (Watts 2002). This is not, of course, how real networks operate: I never see all of my friends at once; particular friends go unseen for months and even years at a time; and, by and large, spending time with one friend comes at the expense of spending time with another.

It is not difficult to see that the contrast between static, concurrent representation and a reality rife with scheduling complexities is consequential for diffusion. For one, it may be that ideas or practices encounter scheduling "bottlenecks" when they reach people with high network centrality, who will take longer than others to query all of their network contacts for their opinions. And if people adopt only when a healthy proportion of their friends seem to have already done so, having many friends who have adopted is not likely to exert much pressure towards conformity if one never encounters them; conversely, having very few friends who have

[6] Network analysts working on disease spreading have been more likely to address the issue of concurrency. An example is Morris and Kretzschmar (1995), who compare simulation models that impose serial monogamy with those that allow for concurrent sexual partners. This does not amount to incorporating scheduling, however, since, except in the serial monogamy model, nothing prevents an actor from having intercourse with several partners on a given day, and indeed a good deal of this goes on in the models. Along similar lines, Moody (2002) considers the effects of network change on diffusion of sexually transmitted diseases, though with similar assumptions.
[7] Burt (1987) describes, and favors, an alternative mechanism, according to which one adopts when those to whom one is "structurally equivalent" – that is, who are similarly tied to other people – have adopted. In this paper I restrict myself to diffusion through direct contact (or through "cohesion," in Burt's terminology), firstly because this is the mechanism more explored by researchers; secondly because it is the one with the most general applicability, including to disease transmission; and thirdly because while the scheduling requirements of diffusion through contact are clear, though normally unstated, the same cannot be said for diffusion through structural equivalence.

adopted *may* induce one to do the same if these are the people with whom one actually spends time.

Before I continue, some terminology: "Diffusion" here refers to the process whereby a practice, piece of information, or disease spreads through a network. "Contagion" refers to a specific occurrence of transmission from one individual to another. I speak of someone as "infected" if they are already in possession of the information or practice or disease, and of "adoption" when an uninfected person becomes infected. I am, of course, shamelessly mixing the terminology of epidemiology with that of the diffusion of innovations, but only in this way can a sufficiently flexible lexicon be compiled without the coining of neologisms.[8]

[FIGURE 3 ABOUT HERE]

The diffusion analysis is diagrammed in Figure 3. It begins with a preference network of 100 actors in which, as already indicated, all ties are constrained to be symmetrical. Diffusion through two types of networks was explored; the algorithms for producing these are described below. The simulation was run 100 times for each combination of conditions (network type and scheduling rule). Each repetition began with the generation of a new network, and the random seeding of the information in a single individual. It ended when everyone in the network had become infected, or after 1,000 rounds, whichever came first.

Network topologies

Much computational research on network diffusion assumes that networks are "random," which normally means that each pair of individuals has some probability of being connected (e.g., Watts 2002). This is useful inasmuch as there are powerful graph-theoretic techniques which tend to make the same assumption, but is clearly not realistic: real networks tend to have particular properties, among them degree centralization – some people are more popular than others (Gould 2002) – as well as a tendency toward clustering, or triadic closure, which means that two people with a common friend tend to be friends with one another (Doreian, Kapuscinski, Krackhardt, and Szczypula 1996). Thus, in what follows I consider diffusion as it occurs in two network topologies: one a random network, in which ties are forged at random, and the other a

[8] Terminological alternatives welcome!

network in which ties are formed in such a way as to generate degree centralization and triadic closure. This will provide us with a basis of comparison for exploring the consequences of scheduling in realistic networks, and for understanding precisely how those consequences come about.

In both sorts of networks, each actor "initiates" five ties, but can be on the receiving end of additional ties, so that the mean degree is ten – five initiated and, on average, five received. In the random (Rand) networks, people forge these ties at random, subject to the constraint that two actors can only have one tie between them, which means that each cannot initiate a tie to the other.[9] The procedure for generating the "centralized-clustered" (C-C) networks is more complicated. Here, actors "arrive" to the network one at a time, and initiate five ties before the next actor "arrives." Ego's first tie is forged according to Barabási and Albert's (1999) principle of "preferential attachment," which calculates the probability of *i* forging this initial tie to *j* as *j*'s proportional share of all existing ties summed across everyone, excluding *i*:

$$\Pr(iXj) = \frac{x_{j+}}{\sum_{k=1, k\neq i}^{g} x_{k+}}. \qquad \text{(Eq. 4)}$$

With ego's initial tie formed in this manner, his or her remaining ties are formed according to a similar calculation of the extent to which individuals *to whom ego is not already tied* share friends with ego:

$$\Pr(iXj) = \frac{\sum_{k=1}^{g} x_{ik}x_{jk}}{\sum_{l=1}^{g}\sum_{k=1}^{g} x_{ik}x_{lk}}. \qquad \text{(Eq. 5)}$$

Equation 6 is applied repeatedly, until ego has forged five ties in all. Note that if *i* finds that there is no one remaining (a) to whom she is not yet tied, and (b) with whom she shares one or more friends – that is, if there is no *j* for whom Equation 6 produces a non-zero probability – she forges her additional ties at random.

---

[9] While it is possible for such a network to be disconnected, something which would create a problem for the diffusion analysis, the odds of this happening are infinitesimally small, and none of the networks generated had this property.

[FIGURE 4 ABOUT HERE]

Sample random and centralized-clustered networks are diagrammed in Figure 4, though here the number of ties initiated by each actor was set at two, to reduce the number of ties so as to make the structural features easier to discern. It is obvious that the centralized-clustered network evidences the two properties of degree centralization, whereby some people have many ties and many others have few, and triadic closure, such that two people are more likely to be friends when they have friends in common; while we also find some relatively central actors, and some closed triads, in the random network, these are at chance levels only. To make the comparison more rigorous, we can calculate the degree of triadic closure using Watts and Strogatz's (1998) "clustering coefficient":

$$C(p) = \frac{\sum_{i=1}^{g} \left(\sum_{j=1}^{g} \sum_{k=1}^{g} x_{ij} x_{ik} x_{jk} \Big/ \sum_{j=1}^{g} \sum_{k=1}^{g} x_{ij} x_{ik} \right)}{g} \qquad \text{(Eq. 6)}$$

where $i \neq k \neq j$. This averages the ego-specific clustering coefficients, each of which is the proportion of pairs of alters liked by ego who also like each other, and has a range of [0, 1]. The mean clustering coefficient for the random graphs, generated as described above, is .093, while that for the centralized-clustered graphs is .431 (in each case, across 100 realizations); the difference is statistically significant (two-sided t-test; $p<.001$). Clearly, the degree of clustering, or triadic closure, in the centralized-clustered networks is much higher than in the random networks, as intended.

[FIGURE 5 ABOUT HERE]

Second, we can compare the degree distributions from the two types of networks for further confirmation that the centralized-clustered networks are, in fact, centralized, meaning that some people have a lot of ties while many others have few. The two degree distributions, aggregating across 100 realizations of each type of network, are presented as histograms in Figure 5. Predictably, the random networks have degree distributions that are approximately normal, with some people having only about five friends (the minimum), some having closer to fifteen, but most falling in the intermediate range of eight to twelve. In contrast, the majority of actors have between five and seven ties in the centralized-clustered networks, while some – a very few – have upwards of *fifty* ties.[10] If we believe Albert and Barabási (1999) and Gould (2002), this is how real networks look, and this particular feature of the centralized-clustered networks will prove especially important in what follows.

First contact contagion

The initial analysis is based on the assumption that contagion occurs on first contact – that is, that ego becomes infected immediately upon encountering an infected alter; imagine a hot piece of gossip or, in the epidemiological application, a highly virulent disease. (Later I consider the effects of "threshold" diffusion, whereby someone adopts only when a particular fraction of their encountered friends already have, as is likely more appropriate for the spread of fads.) I compare mean time to complete diffusion (a) when ties are permitted to operate concurrently, as in traditional diffusion models, (b) when there is scheduling with remediation, and (c) when there is scheduling without remediation, where each analysis was performed on each of the two types of networks, making for six "conditions."

[FIGURE 6 ABOUT HERE]

Figure 6 is a boxplot of rounds until complete saturation for each condition. Several things can be seen here. Most striking, though not surprising, is the delay in diffusion caused by scheduling. Complete diffusion in both random and centralized-clustered networks of 100 actors occurs in about three rounds when actors are permitted to communicate with all of their alters each round – the typical assumption in most network diffusion models. In contrast, complete diffusion in the four scheduling conditions takes, for the most part, between ten and fifteen rounds – this a three-fold to five-fold increase. Clearly, diffusion occurs much more slowly when it requires face-to-face encounters, as may be the case for certain kinds of information and

---

[10] The centralized-clustered distribution is, loosely speaking, "scale-free," though it is not well approximated by the standard power law function $p_k \sim k^{-\gamma}$ (Strogatz 2001), because the decay – the drop-off in frequency of higher degrees – is somewhat too rapid.

gossip, and is certainly the case when we are concerned with the spread of a communicable disease such as SARS. (Compare this to the spread of a computer virus, for which there are no comparable scheduling requirements.)

A second finding is that remediation accelerates diffusion: when actors deliberately service their networks by pursuing elusive alters, the information or disease spreads through the network faster ($p<.001$ for the comparison of random networks with versus without remediation, and $p<.05$ for the comparison of centralized-clustered networks with versus without remediation). What is perhaps more surprising, however, is that the effect is not stronger. It appears that remediation only slightly offsets the effects of scheduling on diffusion time, which, we can thus conclude, are primarily owed to the combination of scheduling constraints and lack of concern with efficient transmission to which I attributed the effect of network topology.

A third finding is that diffusion takes longer in centralized-clustered networks than in random ones: on average, 14.4 versus 10.4 weeks with remediation, and 15.0 versus 11.9 weeks without it. (Both two-sided t-tests resulted in $p < .001$.) It is possible that this difference has at its root the corresponding difference in geodesics, as reflected in time to complete diffusion when no scheduling constraints apply (in which circumstance geodesics primarily determine diffusion time). This difference, however, is very small, to the point that it is not visibly discernable in Figure 6: without scheduling, complete diffusion takes an average of 3.03 rounds in the random networks, and 3.21 rounds in the centralized-clustered networks. This is smaller than the differences we observe when scheduling constraints apply, both in absolute terms, *and* in relative terms.[11] I take this as evidence that the difference in diffusion times between the two sorts of networks when scheduling applies is not simply the result of the initial difference in geodesics as amplified by scheduling constraints, but is due to something else.

[FIGURE 7 ABOUT HERE]

This “something else” can be gleaned from Figure 7, which graphs the round on which an actor was infected against its degree, for 2,000 actors, each drawn at random from a unique repetition of the simulation, half of them from random networks and half from centralized-

[11] That is, 14.4 – 10.4 = 4 is a smaller percentage of 14.4, and 15 – 11.9 = 3.1 is a smaller percentage of 15, than 3.21 – 3.03 = .18 is of 3.2.

clustered networks. (Actors schedule remediatively.) It appears that the longer diffusion times in centralized-clustered networks are due to central actors who bridge between network regions which are, if not otherwise disconnected, at least far-removed from one another. Because these actors have so many alters to answer to, each takes longer than others to become infected once one of its neighbors is. While, once infected, a centralized actor is able to quickly infect others, the initial delay is enough to slow diffusion time relative to what it is in a random graph, because non-central actors on the opposite end of the network from the initial seed take a long time to get infected, having to wait either for a central actor to become infected and to visit them, or for the disease to find its way to them through a more circuitous route. In Figure 7, this is suggested by the triangular configuration of points for the centralized-clustered networks, and in particular by the low-degree, late-adopters that make up the lower-right corner.

Threshold contagion

I have assumed in the foregoing that an actor acquires the information or virus in question immediately upon encountering someone who already has it. While, as noted already, this makes sense in the case of a particularly juicy bit of gossip or a particularly contagious disease, it makes less sense if our concern is with the diffusion of a practice or fashion, where actor is more likely to adopt when some proportion of his or her network alters have – this being that person's "threshold." What a threshold means in scheduling terms is different than what it means in a concurrency model, however. If, as I have been assuming, people are only influenced by friends whom they encounter, then the relevant calculation is not of the proportion of one's total friends who are infected, but the proportion of those friends *one has encountered over some finite period*. We thus have two parameters to consider, in addition to scheduling rule and network type: one's threshold, or the proportion of friends one has recently encountered who are infected before one decides to do adopt, and one's "memory," or the number of rounds one thinks back in deciding whether or not to adopt – where this is what defines "recent."[12]

I assume, for the sake of simplicity, homogeneity in both memory and threshold, which is to say that these are fixed at a single level for all members of the network, and then manipulated

[12] An alternative would be to define memory as a decay function of the recent past. Another option, were we less interested in memory as a parameter, would be to calculate an actor's probability of adopting as a function of the proportion of his or her friends who have already adopted.

to allow us to examine their effects on diffusion. The results are presented in Figure 8. For each condition – combination of network type and scheduling rule – two graphs are shown. The first indicates the proportion of repetitions at each memory-threshold level that resulted in complete diffusion *within 1000 rounds*. (As indicated earlier, the simulation terminated if complete diffusion had not occurred by the 1,000th round.) The second indicates the number of rounds to complete diffusion, up to the maximum of 1000. The graphs in the right column thus understate the time to complete diffusion when this was not achieved within the allotted period.

[FIGURE 8 ABOUT HERE]

When thresholds are very small and/or memories are very short, diffusion takes about as long as it did in the first-contact model. This is not surprising, since under these conditions, one exposure to an "infected" individual is sufficient for contagion, so that threshold contagion becomes, for all intents and purposes, first-contact contagion. If, in particular, memory = 2, a single encounter is enough to meet or surpass a threshold of up to .5, the highest considered here. The same applies if the threshold = .1 so long as memory does not exceed ten, as well as if threshold = .2 and memory = 4. I will refer to any threshold that is satisfied by a single encounter with an infected person to be a "trivial" threshold. Not only does diffusion occur at about the same pace when thresholds are trivial as in the actual first-contact model, the effects of remediation and network structure observed in Figure 6 also carry through, with remediation again expediting diffusion, and the centralized-clustered network again retarding it. (Due to the range of values on the vertical axis, these effects, modest in comparison to what follows, are not visible in Figure 8.)

Diffusion time begins mounting as soon as we venture outside of this area. And yet, so long as memory remains under eight rounds, and/or threshold remains under .4, complete diffusion almost always occurs within 1,000 rounds, and usually in well under half that. While this does indeed amount to a large increase in diffusion time, before dwelling upon this it is important to understand that diffusion would *never* occur at higher threshold levels were it not for scheduling. The reason is that if people were continually monitoring all of their friends all of the time, no one with several friends – and I assume here that everyone has at least five – would ever judge that a significant proportion of her friends had adopted when, particularly when, at the very beginning of the process, only the "seed" is in possession of the innovation. What scheduling does is repeatedly bring a person into the presence of those few friends – and there may only be one – who have adopted, with the result that she incorrectly *infers* that a larger fraction of her friends have adopted than actually have, and thus adopts. The implication is that some innovations, at least, initially take hold in a population because some people "misread" their ego-networks, jumping to conclusions based on scheduling contingencies that they only imperfectly understand.

Different mechanisms are responsible for diffusion when memory is in the medium range and threshold is low, on the one hand, and when threshold is in the medium range and memory is short, on the other. When threshold levels are moderate (here, .2 or .3), and memory is long (8-10 rounds), contagion can occur because people cast a wide net and adopt if a small subset of those captured are found to have already adopted. We might refer to this as contagion through "high sensitivity," and we can expect it to occur when people are eager to innovate for the sake of innovation, as in managerial circles when the consequences of an innovation are of less concern than the appearance of remaining on the "cutting edge." In contrast, when memory length is moderate (4-6 rounds) and threshold is large (.4-.5), contagion can occur when, as just explained, chance alone lands someone in the company of that small subset of his or her friends who have already adopted, on the basis of which he or she infers, perhaps incorrectly, that most of his or her friends (i.e., not just those recently encountered) have done so, and thus adopts. This is diffusion through small-N sampling, or "stochastic contagion," and we might expect to find it occurring among adolescents concerned with keeping up with quickly-changing fashions.

This sets us up to understand how diffusion times can be so long when memories and long and thresholds are high – and indeed, why sometimes diffusion appears not to happen at all, at least within the allotted time. The problem with stochastic contagion is that it requires short memories: the further back into the past people look, the more encounters, and thus alters, they are able to consider, the less likely they are to jump to incorrect conclusions about the fraction of their friends who have adopted. On the other hand, contagion through sensitivity requires that people have low thresholds; otherwise, a long memory just means that people are better able to enumerate counter-examples to those one or two friends who have already adopted. When memory is long *and* threshold is high, people spend a long time examining their friends before they consider adopting, and then only adopt if a sizeable portion of recently-encountered friends have – a recipe, as it turns out, for stalled diffusion. This is what we might expect of people wary of innovation, such as social scientists reluctant to adopt a new statistical technique for fear that reviewers will not understand it, or intelligence officials more concerned with accountability according to established standards than with maximum effectiveness or efficiency.

The same combination of moderate-to-high threshold and medium-to-long memory that retards diffusion has two further consequences. I begin with the simpler of these, though it is the less stark. While remediation expedites diffusion when contagion occurs upon first contact, and when thresholds are trivial, remediation actually *slows* threshold-based diffusion when memory is medium to long and threshold is moderate to high.[13] In general, remediation reduces the likelihood that an actor will, by chance alone, encounter only that subset of his or her friends who have already adopted.. This is consequential when memory is medium to long and threshold is moderate to high because, under these circumstances and early in the diffusion process, a person has to repeatedly encounter only that friend who has already adopted for his or her threshold to be reached, and this is exactly what remediation militates against. While we might, consequently, expect the remediation that professionals engage in to accelerate the diffusion of a rumor, it may well retard the spread of an uncertain new practice, the adoption of which one is reluctant to take on absent the belief that others (note the plural) have already done so.

The other thing that happens when threshold is moderate-to-high and memory is medium-to-long is that the effect of network structure is reversed, such that diffusion takes longer in random networks than in centralized-clustered networks – *much* longer when threshold and memory approach their upper values.[14] (Compare Figures 8b and 8f, and 8d and 8h.) That is, while diffusion occurs faster in random networks than in more realistically structured ones when people adopt upon first contact, or when, more generally, a single encounter with someone already infected is sufficient for contagion to occur, it occurs faster in realistic networks than in

[13] Determining statistical significance in this case is complicated, firstly because any two points on the planes in Figures 7b and d, and f and h, could be compared, and secondly of the right censuring caused by setting the maximum number of rounds at 1000. However, to take one set of memory-threshold coordinates, when threshold = .3 and memory = 2, the difference between mean rounds in Figures 7b and d is significant at the $p < .001$ level, and between mean rounds in Figure 8f and h at the $p < .01$ level.

[14] This effect, and that related to remediation described earlier, does not necessarily take hold at the instant that thresholds become non-trivial, which is to say that there are low-to-intermediate levels of memory and threshold at which the first-contact regime still faintly applies. As exact the transition point depends on a number of considerations – in particular, whether we are concerned with network structure or remediation – and as it is not sharp in any case (i.e., it does not amount to an abrupt “phase transition”), I do not dwell upon it here.

random ones when contagion occurs contingent upon a significant fractional threshold being reached.

[FIGURE 9 ABOUT HERE]

It turns out that the difference between the two sorts of networks is entirely due to what transpires in the opening rounds of the simulation. This is conveyed by Figure 9, which graphs, for each combination of network type and remediation or its absence, the number of rounds passing between successive instances of contagion, aggregating across 100 repetitions, where threshold is set at .5 and memory to 6. Thus, in random networks with remediation the seed takes an average of 129 rounds to infect just one person, following which the second, third, and fourth individuals are infected – not necessarily through the seed's direct influence – at average intervals of forty-seven, twenty-seven, and eighteen rounds.

[FIGURE 10 ABOUT HERE]

What we see here is that the rates of diffusion under the four conditions quickly converge, such that by the tenth round they are essentially indistinguishable. What is happening in these opening rounds? The answer is in Figure 10, which, like Figure 7, graphs an actor's degree against the round in which it was infected. From the insert it is evident that the earliest adopters are those with especially low degree, and who are as a consequence most susceptible to the influence of a single infected individual. (The gap between the seeds, who are infected at round 0, and the first subsequent adopters is a product of memory length.) There is no comparable pattern for the corresponding graph for clustering (not shown): some early adopters have highly clustered networks, while others have very unclustered ones. Further, note that the earliest adopters from the random networks are those whose degree best approximates what is modal in the centralized-clustered networks. But lacking many such people, diffusion in random networks takes considerably longer to get off the ground.

Once diffusion does get off the ground, the pattern of who adopts when it similar to what we saw in connection with first-contact diffusion in Figure 7. Once again, central actors in the centralized-clustered networks adopt at an intermediate stage, yielding to the influence of the more peripheral early adopters. Then these central actors assist in the further transmission of the practice or idea to peripheral actors on the other side of the network. [Maybe add something here about the longer tail in Figure 10 – this may suggest that central actors don't actually exercise very effective influence over peripheral ones.]

The implication is that topological differences are most consequential in the earliest stages of the diffusion process, at least with respect to the question of how long complete diffusion takes. And, judging from Figure 9, this is also where remediation has the greatest impact, by temporarily inhibiting diffusion in the crucial opening rounds. [Needs development.]

Discussion

The starting point of this paper was the gap between network analysts' representation of social networks as configurations of currently existing and active ties, and the reality of episodic encounters, scheduling difficulties, and prolonged periods of tie inactivity. By means of a computer simulation of network scheduling, I have demonstrated that this disjuncture is consequential, at least for the phenomenon of diffusion. This consequentiality, however, is not at all straightforward. The main findings, including the two contagion mechanisms, are summarized in Figure 11.

[FIGURE 11 ABOUT HERE]

I submit that these findings are likely robust to the precise details of the scheduling algorithm, since they follow more from the undeniable constraints upon which it is based – such as the constraint that a person cannot be in two places at once – than upon, say, the precise implementation of impatience. Still, here is far more to scheduling than is captured in this simple simulation, and some of it may be consequential for not only diffusion, but also for other outcomes that it might be used to explore. Among the scheduling considerations lacking from this model are: that people can schedule encounters far in advance, committing to future encounters with people who cannot be immediately accommodated; that people can meet in groups of larger than two, though there are good reasons why people prefer one-on-one encounters; that for a range of reasons, people block off time for solitary work or reflection, sometimes in response to temporal rhythms of an entirely different kind (e.g., related to organizational routines); that people can affix different degrees of urgency to their requests for encounters; that people hold some friends in higher regard than others, and may positively dislike some people; and that people can employ heterogeneous scheduling strategies, whereas here I have assumed homogeneous scheduling "cultures."

Each of these variations is well worth pursuing, in connection with the outcomes explored above as well as others, such as network evolution, social integration/involvement, and organizational decision-making. I wish to suggest, however, that this paper points to more than just further extensions of the model. While my specific concern in this paper has been with scheduling, I take this to be one aspect of a more general set of issues related to what I call "enactment dynamics." As here, the study of enactment dynamics begins with a distinction between historically-anchored, pre-existing relationships, on the one hand, and the various ways in which these can be concretely activated, or "enacted," on the other. Enactment need not entail simply having lunch for a friend, though this is the sort of enactment assumed in this paper. It may also entail various sorts of actions towards third parties (e.g., one's friend's enemy), with or without deliberate coordination with the person to whom the tie extends. The important thing is that such action would not have occurred absent the tie in question, though in defining "tie enactment" we might also add the requirements that the action not undermine that tie, and maybe also that the action be recognizable to third parties as something one is doing by virtue of the tie.

The study of enactment dynamics stands to extend the reach of network analysis, firstly by encouraging us to view a wider range of actions as products of our network positions, and secondly by alerting us to the myriad obstacles that stand between a pre-existing network and any consequences following from it in a particular setting, here exemplified by scheduling constraints that may keep friends from encountering one another for long periods. While such obstacles might seem at first to pose a threat to network analysis, by blocking network effects, an understanding of these constraints will better equip us to understand how networks operate through and around them, as well as to anticipate the situational junctures at which particular networks seem to matter not at all.

Conclusion

The premise of this paper is that standard network representations are fundamentally misleading about the nature of networks, implying as they do that a person's network ties are continuously in operation, though in reality we know that ties can be activated and deactivated, enacted and suppressed, and that the activation of one tie often has as its consequence the suppression of others. Central to the translation of a "background" preference network into concrete patterns of interaction are principles of scheduling, by means of which people seek to balance, not always successfully, the competing demands incumbent upon their network positions. By means of a scheduling simulation that is built, to the extent that a simple simulation can be, upon a realistic phenomenology, I have sought to show that the scheduling constraints that interpolate between network representations and network realities are consequential for at least one outcome of concern to network analysts, social scientists generally, and recently, the wider public: diffusion. The effects of scheduling were found to depend upon several factors, including the behavioral (adoption) rule scheduled encounters inform, the network structure scheduling operates upon, and how actors themselves respond to scheduling outcomes, in terms of whether or not they strive to remedy past imperfections. I concluded with a call for further research on "enactment dynamics" more generally, encompassing not only network scheduling but all of the various circumstances and constraints that activate and suppress network ties, and that thus decide when, and how, pre-existing relations translate into events in the world.

Works cited

Albert, Réka and Albert-László Barabási. 2002. "Statistical Mechanics of Complex Networks." *Review of Modern Physics* 74:47-97.

Axelrod, Robert. 1997. "Advancing the Art of Simulation in the Social Sciences." Pp. 21-40 in *Simulating Social Phenomena*, edited by R. Conte, R. Hegselmann, and P. Terna. Berlin: Springer.

Barabási, Albert-László and Réka Albert. 1999. "Emerging of Scaling in Random Networks." *Science* 286:509-12.

Burt, Ronald S. 1987. "Social Contagion and Innovation: Cohesion versus Structural Equivalence." *American Journal of Sociology*:1287-335.

—. 1992. *Structural Holes*. Cambridge, MA: Cambridge University Press.

Cohen, Sheldon, William J. Doyle, David P. Skoner, Bruce S. Rabin, and Jack M. Gwaltney. 1997. "Social Ties and Susceptibility to the Common Cold." *Journal of the American Medical Association* 277:1940-44.

Coleman, James, Elihu Katz, and Menzel. 1967. *Medical Innovation*.

Crystal, David and Hilary Crystal. 2000. *Words on Words: Quotations About Language and Languages*. Chicago: University of Chicago.

Doreian, Patrick, Roman Kapuscinski, David Krackhardt, and Janusz Szczypula. 1996. "A Brief History of Balance Through Time." *Journal of Mathematical Sociology* 21:113-31.

Epstein, Joshua M. and Robert Axtell. 1996. *Growing Artificial Societies*. Washington, D.C.: Brookings Institute.

Galaskiewicz, Joseph and Ronald S. Burt. 1991. "Interorganization Contagion in Corporate Philanthropy." *Administrative Science Quarterly* 1991:88-105.

Gilbert, Nigel and Klaus Troitzsch. 1999. *Simulation for Social Scientists*: Open University Press.

Goffman, Erving. 1967. *Interaction Ritual*. New York: Pantheon.

Gould, Roger V. 2002. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." *American Journal of Sociology* 107:1143-78.

Leifer, Eric M. 1990. "Enacting Networks: The Feasibility of Fairness." *Social Networks* 12:1-25.

Macy, Michael W. and Robert Willer. 2001. "From Factors to Actors: Computational Sociology and Agent-Based Modeling." *Annual Review of Sociology* 28:143-66.

McPherson, Miller J., Pamela A. Popielarz, and Sonja Drobnic. 1992. "Social Networks and Organizational Dynamics." *American Sociological Review* 57:153-171.

Moody, James. 2002. "The Importance of Relationship Timing for Diffusion." *Social Forces* 81:25-56.

Morris, Martina and Mirjam Kretzschmar. 1995. "Concurrent Partnerships and Transmission Dynamics in Networks." *Social Networks* 17:299-318.

Pattison, Philippa and Stanley Wasserman. 1999. "Logit Models and Logistic Regressions for Social Networks, II. Multivariate Relations." *British Journal of Mathematical and Statistical Psychology* 52:169-94.

Podolny, Joel M. and James N. Baron. 1997. "Resources and Relationships: Social Networks and Mobility in the Workplace." *American Sociological Review* 62:673-93.

Skyrms, Brian and Robin Pemantle. 2000. "A Dynamic Model of Social Network Formation." *PNAS* 97:9340-46.
Strogatz, Steve. 2001. "Exploring Complex Networks." *Nature* 410:268-76.
Wasserman, Stanley and Katherine Faust. 1994. *Social Network Analysis*. Cambridge: Cambridge University Press.
Watts, Duncan J. 2002. "A Simple Model of Global Cascades in Random Networks." *PNAS* 99:5766-71.
Watts, Duncan J. and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small-world' Networks." *Nature* 393:440-2.
Winship, Christopher. 1978. "The Allocation of Time Among Individuals." Pp. 75-100 in *Sociological Methodology*, edited by K. F. Schuessler. San Francisco: Jossey-Bass.
—. 1992. "Social Relations and Time." Unpublished.

Figure 1. Scheduling algorithm core

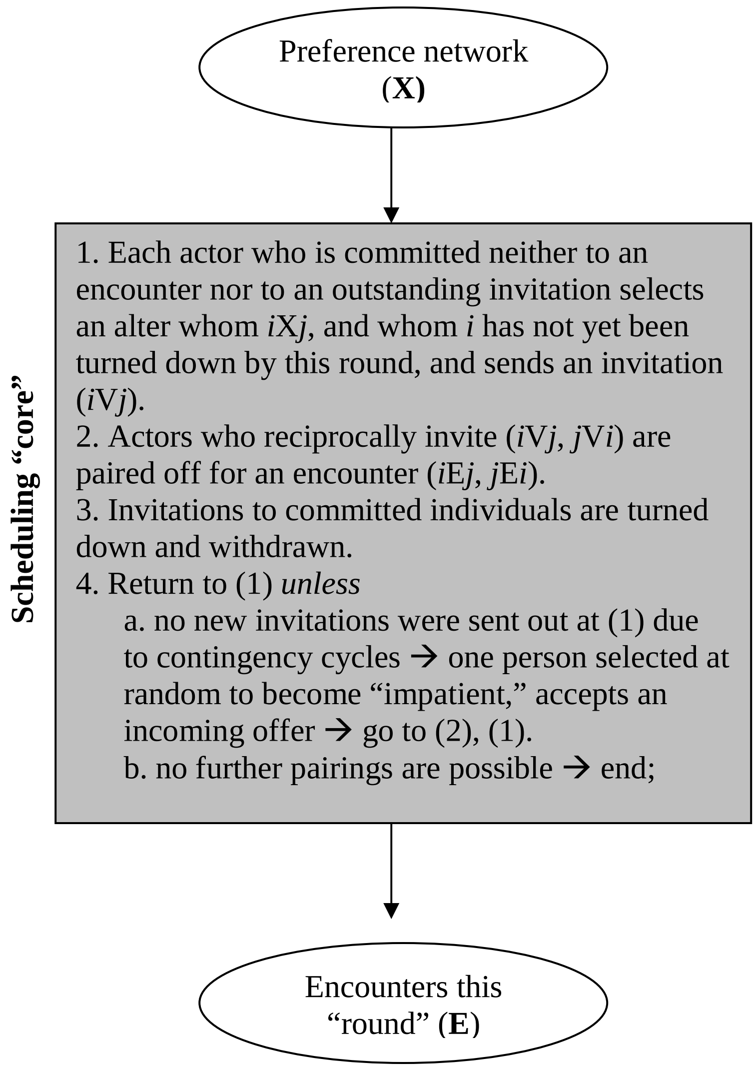

Figure 2. Invitation contingency cycles (arrows are invitations)

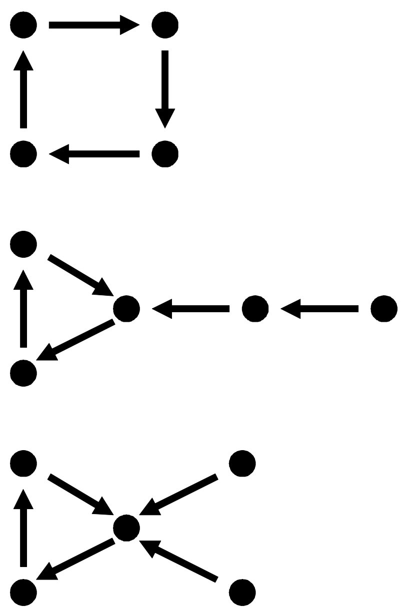

Figure 3. Diffusion analysis

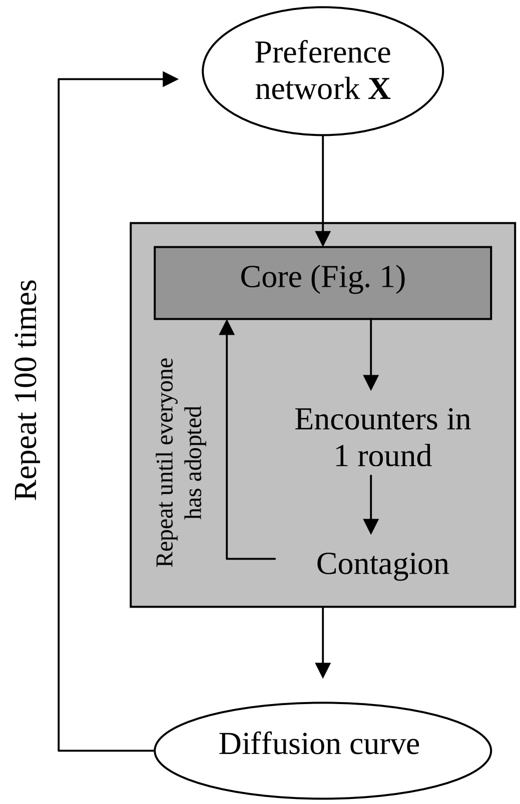

Figure 4. Sample networks, minimum degree = 2

a. Random

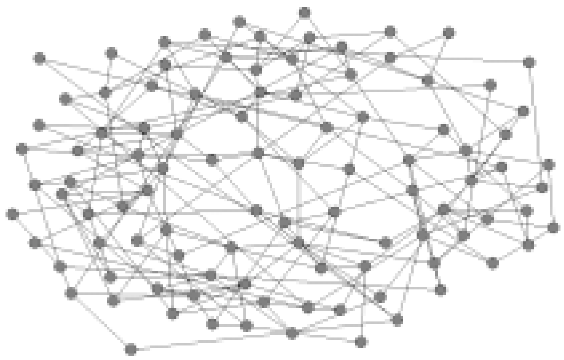

b. Centralized-clustered

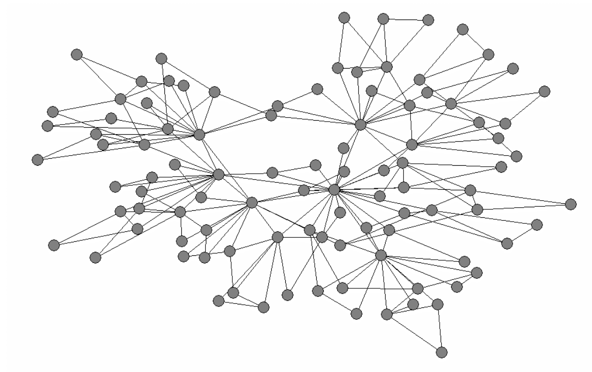

Figure 5. Degree distributions, summing over 100 networks, minimum degree = 5

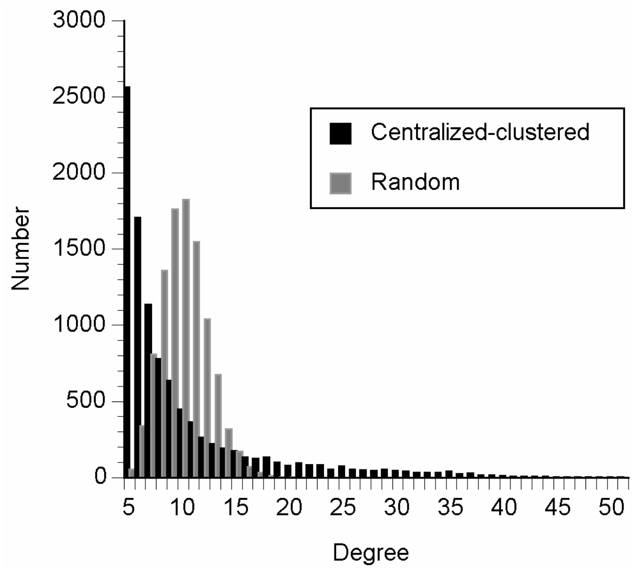

Figure 6. Boxplot for rounds to complete diffusion by network and rule combination

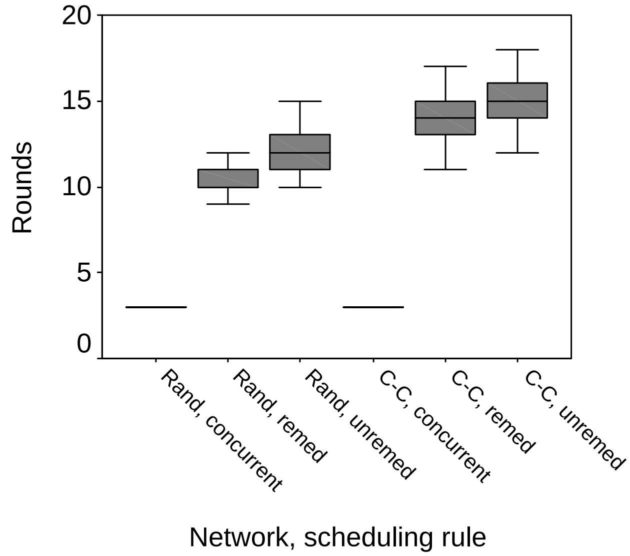

Figure 7. Time to infection by degree (given remediation), first-contact model

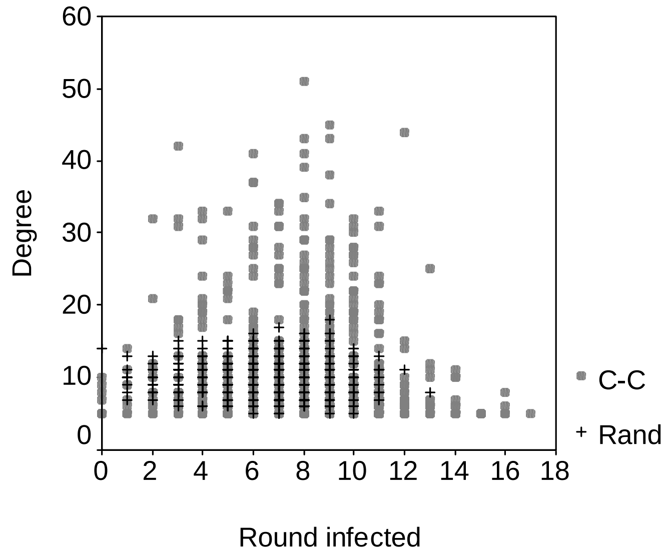

Figure 8. Rounds to complete diffusion, and proportion completely diffused within 1000 rounds, by network and rule combination

	Proportion that completely diffused within 1000 rounds	Mean rounds to complete diffusion (max=1000)
Rand, remed	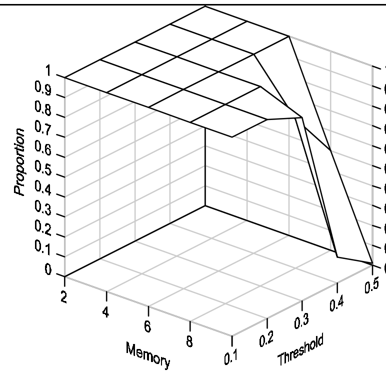 a.	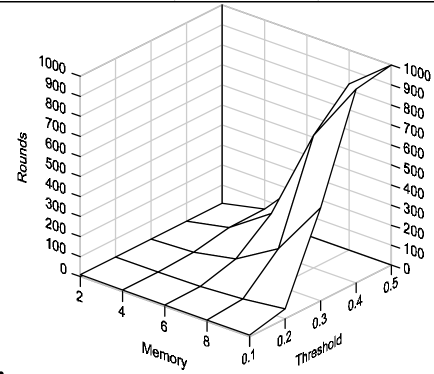 b.
Rand, unremed	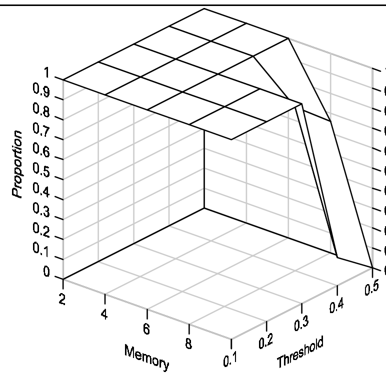 c.	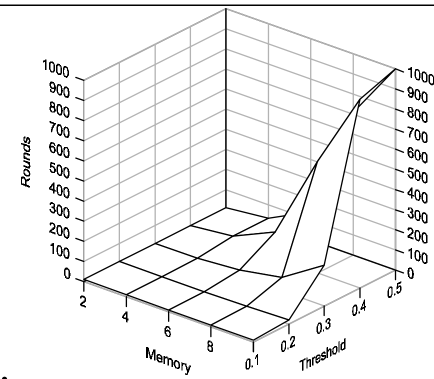 d.
C-C, remed	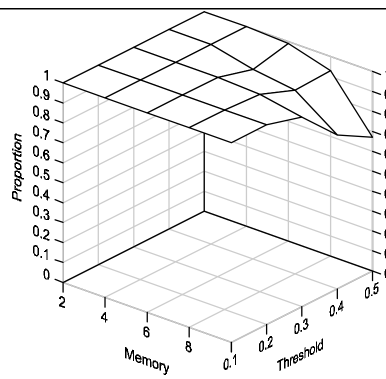 e.	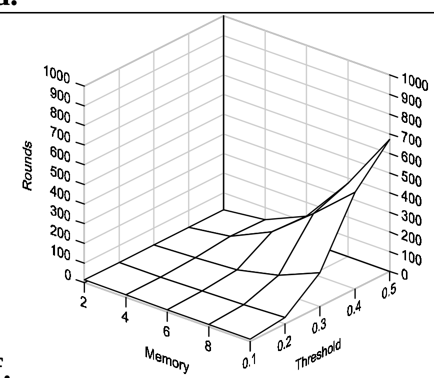 f.
C-C, unremed	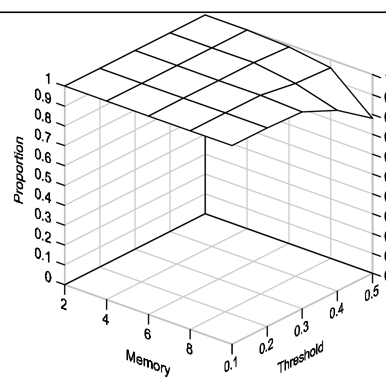 g.	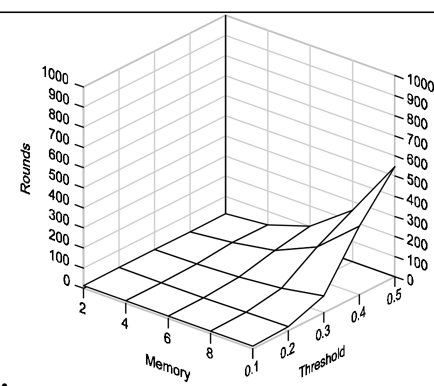 h.

Figure 9. Rounds to next contagion, given last

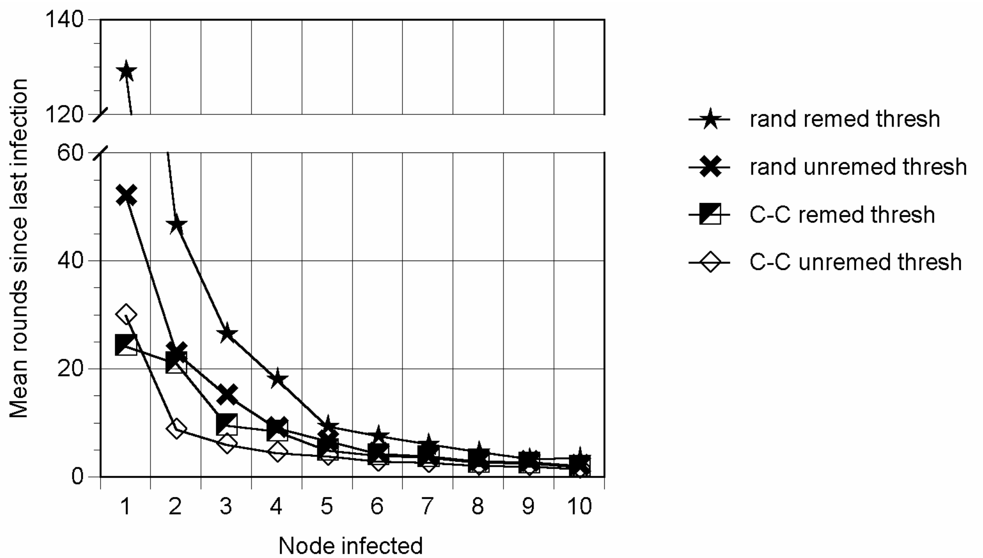

Figure 10. Round infected by degree, for threshold = .5, memory = 6, and remediation (1,000 actors, each drawn from a separate repetition)

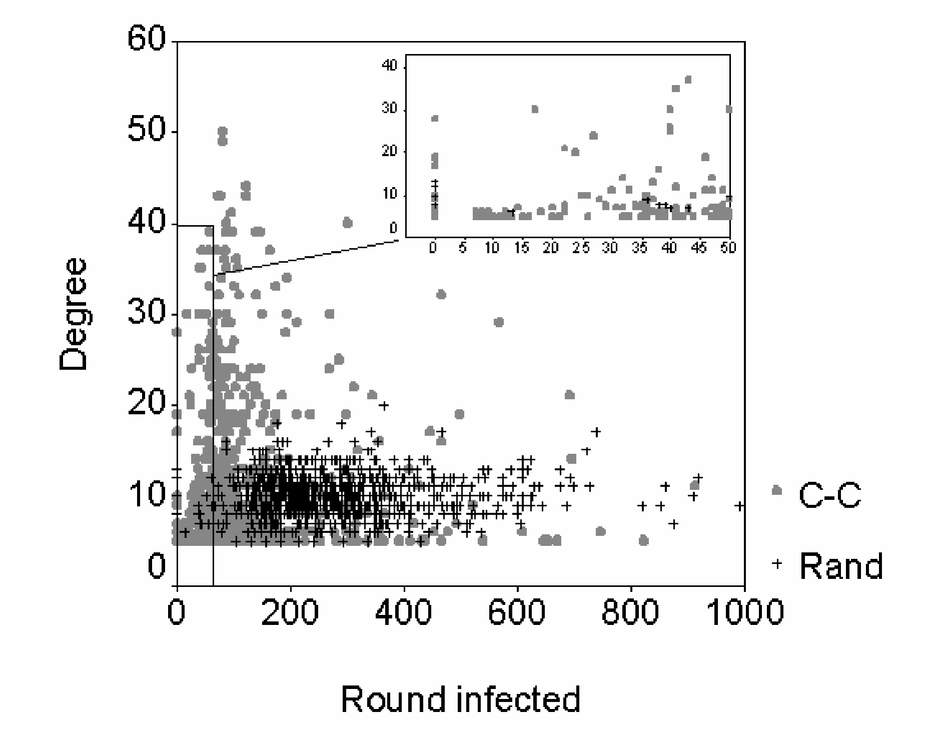

Figure 11. Summary of main findings

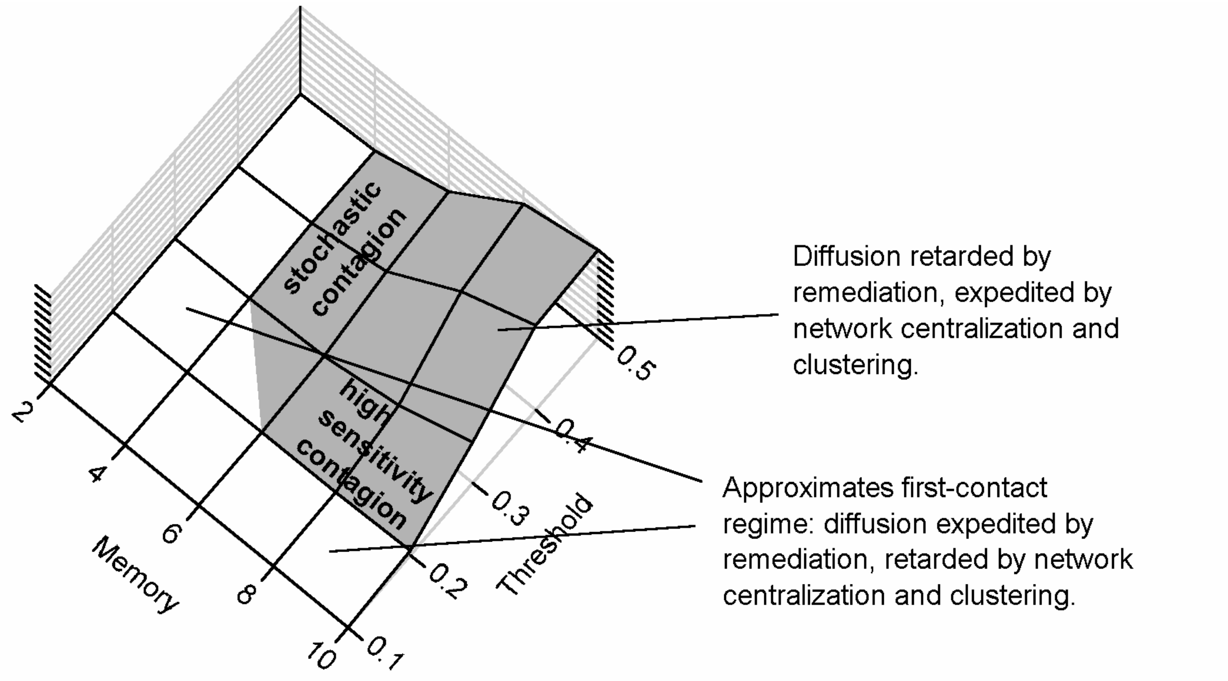

TO DO: Change language: to diseases in the 1st-contact model, and practices in the threshold model. And explain earlier that the two models are meant for these different sorts of phenomena. Future work: parameterize clustering and degree centralization. Here I can’t because it’s hard and maybe, at this stage, impossible to only vary one feature of a network; clustering and centralization, in particular, tend to go together.