Modeling Community Resilience for a Post-Epidemic Society

Shaun Michel[1] and Karine Megerdoomian[1]

[1] George Mason University, Department of Computational Social Science, Research 1, 4400 University Drive, Fairfax, VA 22030, USA
{smichel3, kmegerdo}@gmu.edu

Abstract. The 2014 Ebola outbreak in West Africa once again reminded the world of the fatal risks of exposure to deadly disease. Commonly evaluated by the number of fatalities estimated during an outbreak, epidemics also have lasting consequences for survivors in the form of community breakdown. Although agent-based models are frequently used to consider the susceptibility of agents to disease and to predict the evolution of epidemics, they rarely attempt to model the interplay of social networks, spatial awareness, and exposure risks. Our response is to construct a model in which agents are instantiated within a social network that influences their movement decisions alongside individually perceived vulnerabilities to exposure-by-contact. We monitor macroscopic behavioral trends and examine community breakdown resulting from fatalities. The model provides an important contribution to modeling social science by exploring individual response to an emerging epidemic and community-level outcomes as a result of those responses.

1 Introduction

According to the US Center for Disease Control (CDC), the 2014 Ebola epidemic was the largest in history, affecting multiple countries in West Africa and resulting in tens of thousands of deaths [1]. The United Nations has predicted long-term socio-economic impact in West Africa due to the epidemic. The severity of the 2014 Ebola epidemic has highlighted the importance of preparedness and mitigation strategies, and strengthening community resilience has recently emerged as an essential component of disease outbreak and diffusion mitigation [2].

Epidemics propagate through a social network and can break up communities by causing losses in families and effectively disconnecting the nodes in the network. However, the behavior of the people after the disease and the presence of societal institutions can allow the community to reorganize and rebuild a networked community after the disaster has passed. We built an agent-based model (ABM) to explore structural characteristics and individual beliefs that influence the resilience of a community to an emergent threat such as the Ebola epidemic.

In this paper, we look at social factors beyond fatality which can have long-term consequences on entire communities. We consider that social capital, namely the initial social structure as well as the decisions that individuals make, can impact a community's resilience against disastrous events such as an epidemic. We model a stylized

networked community affected by disease propagation causing fatalities, in order to study the adaptive capacity of the network in a post-disaster environment.

2 Resilience

Resilience, broadly defined as the capacity to resist and recover from loss is considered a desirable property of natural and human systems. The term was first coined by Holling in studying the ecological system's response to natural disasters, who described it as a ''measure of the persistence of systems and their ability to absorb change and disturbance and still maintain the same relationships between populations or state variables [3].'' According to this definition, a resilient system is not necessarily a stable one since a very stable system would not fluctuate greatly and will return to normal quickly, while a highly resilient system may become quite unstable when perturbed, undergoing significant fluctuation, but having the ability to recover and return to a state of equilibrium [4].

Since the seminal work by Holling [3], alternative definitions have been provided based on different system properties and distinct domains of application. Pimm [5], for instance, defines resilience as the speed with which a system returns to its original state following a perturbation. The concept has also gained ground in social science, where it is applied to describe the behavioral response of communities, institutions and economies. Timmerman [6] defines resilience as the ability of human communities to withstand external shocks or perturbations to their infrastructure such as environmental variability or social, economic or political upheaval, and to recover from such perturbations (also see [7]).

Kimhi and Shamai identify three main directions of analysis in the literature in describing community resilience: resistance, recovery and creativity, in which (1) resistance relates to the ability of a community to withstand a disturbance and its consequences, and can be understood in terms of the degree of disruption that can be accommodated without social entity undergoing long-term change; (2) recovery relates to an entity's ability to pull through the disturbance, and can be understood in terms of the time taken for an entity to recover from a disruption; and (3) creativity which addresses the ability of a social system to maintain a constant process of creating and recreating, so that the community not only responds to adversity, but in doing so, reaches a higher level of functioning [8].

In this paper, we define resilience as a property that reflects the capacity of a system to cope with disturbance and reorganize while undergoing change to maintain structure and functioning [9]. One important factor that emerges from the literature on community resilience is the need for the community to return to a system that more or less maintained the original structure or if it has been modified, it has been able to preserve the original functioning of the societal system.

While Adger focuses on community resilience in terms of resource dependency [5], Breton [10], Adger [11] and Aldrich [12] each argue for the essential role of social capital, consisting of an adequate services infrastructure and networks of people and local voluntary associations through which members of the community can be mobilized for action. Adger [13] claims that adaptation processes to natural disaster involve

the interdependence of agents through their relationships with each other, with the institutions in which they reside, and with the resource base on which they depend. Based on several case studies, Aldrich [12] makes a stronger case for the role of social capital and claims that the resilience of a community and its ability to recover after a disaster crucially depend on the internal social networks: "higher levels of social capital —more than factors such as greater economic resources, assistance from the government or outside agencies, and low levels of damage—facilitate recovery and help survivors coordinate for more effective reconstruction." According to Aldrich, even highly damaged communities with low income and lack of access to governmental resources and outside aid can recover if they include dense social networks and tighter bonds with relatives, neighbors, and extralocal acquaintances.

Aldrich [12, 14] also explains variation in recovery among various post-disaster societies, by arguing that social capital should be developed and reinforced at all levels (bonding, bridging, and linking social capital) for best results. Bonding social capital refers to social bonds within and between family, kin, and ethnic group members. Bridging social capital connects members of the group or network to extra-local networks, crossing ethnic, racial, religious groups. Bridging activities and organizations can bring together individuals from different neighborhoods, racial and ethnic identities, and language groups. Finally, linking social capital is made up of "networks of trusting relationships between people who are interacting across explicit, formal or institutionalized power or authority gradients in society". Where bonding social capital facilitates cooperation among members of the same family or neighborhood, bridging capital involves linkage to external assets and generates broader identities. In contrast with bonding and bridging social capital, which primarily connect individuals of the same status, linking social capital takes into account vertical distance as well. By contrasting the recovery of various social areas after a disaster, Aldrich argues that areas with both bonding and linking resources fare better than villages and neighborhoods which have only bonding social capital.

3 Modeling Resilience to an Ebola Epidemic

3.1 Background

The computational work on the Ebola epidemic has focused mainly on approximating and projecting the evolution of the spread of the disease [15,16]. It has long been recognized that the structure of a social network plays an important role in the dynamics of disease propagation. Models of the Ebola epidemic explore network topology in how quickly the disease can diffuse through a social network, what trajectory it might take, and how resilient the network would be to the deletion of vertices. In general, a small world network tends to facilitate disease propagation, while scale free (preferential) networks are found to be resilient to the random deletion of nodes but susceptible to the targeted deletion of hubs. This also has consequences for policy development and immunization strategies, as health professionals can target specific nodes for better results. Nevertheless, the research has not concentrated on the determination of the level of post-epidemic resilience provided by distinct initial network types.

In addition, recent sociological and health policy studies have pointed to the importance of cultural beliefs and behavioral responses in affecting the propagation of the Ebola outbreak in West Africa. Some of these local beliefs include denial of the virus' existence, mistrust of authorities, fear of going to the hospital, and traditional beliefs in disease causation [16,18]. The strength of individuals' denial when dealing with the Ebola outbreak is discussed in [19], while [20] argues that the impact of fear, denial, and stigmatization should not be neglected in developing health policies in African countries. Although these individual beliefs have not been included in the existing models of the 2014 Ebola epidemic, the media coverage offered insightful articles on the role of such cultural elements and individuals' beliefs: [21] focuses on the role of the Ebola epidemic in destroying families in a country where extended families are the most important institution, and wonders about the long-term impact on keeping societies together. This New York Times article describes the fate of an extended family in Liberia that pulled together to take care of infected family members. The care-taking family members were consequently infected with the virus, resulting in a large number of deaths in the family and in survivors being ostracized by the community. The article points to a few interesting elements: the role of denial of the epidemic —at a societal level, but also at the individual level —and the role of the central figure in the family in propagating that denial. As the article describes in detail, the denial by the matriarch of the family—and her active deception of other members of the family concerning the severity of the disease—plays a crucial role in the large number of fatalities suffered by this Liberian family. By focusing on the role a woman plays in recovering from the pandemic, an NPR article [22], describes villagers who have lost family members to Ebola but are showing resilience through their relationships with their neighbors.

These news articles on the 2014 epidemic and the results obtained from sociological and anthropological studies from previous Ebola epidemics point to individual attributes that, combined with a social network structure, can have an effect on community resilience during an epidemic and shortly after it, by either jeopardizing the population and the social capital through denial and risk-taking or by rebuilding the community by (re)establishing social relations with neighbors and other members of the society. In this paper, we study how the societal network structure at the onset of the epidemic combined with individuals' cognitive biases and strategies for making decisions shape the post-epidemic social environment. In particular, we propose to build an ABM to explore the role of the social network and the impact of individual beliefs, such as fear and denial, on community resilience following an epidemic.

In developing this model, we are interested in detecting the contributions towards community resilience of both structural and cognitive properties of social capital. Structural social capital refers to the networks and social structures in which people are embedded, while cognitive social capital refers to contents, such as trust, shared beliefs, norms of obligation and reciprocity [23]. The ABM developed identifies the effect of structural factors such as the initial community relations and network structure and structural position of the casualties due to the disease. In addition, the model takes into account cognitive factors by exploring the impact of individual beliefs such as denial and fear. In this implementation of the model, however, we do not distinguish between levels of social capital, such as bonding, bridging and linking ties. Thus, all nodes are treated equally as members of the community and we do not distinguish civic and faith-based organizations (e.g., churches) or the political structure.

3.2 Measuring Resilience

The determination of a valid measure for community resilience is still being debated in the field. Frankenberger et al [24] proposes a measurement framework for community resilience to model the dynamics of resilience capacities (capacity for collective action) in relation to key well-being outcome indicators (e.g., food security, health/nutrition, access to services, social capital) and shocks and stressors. Cutter et al [25] identify comprehensive resilience indicators by treating resilience as a dynamic process dependent on antecedent conditions, the disaster's severity, time between hazard events, and influences from exogenous factors. For the purposes of this paper, as we are mainly focusing on the role of social capital in determining community resilience, we study factors related to the social network and how it changes as a result of the shock (i.e., disease) applied to the community. We therefore capture the system's resilience based on the resulting fragmentation of the network, measured by the final number of clusters obtained.

The paper addresses the following research questions:

1. How does the original structure of the network affect resilience of the community?
We hypothesize that, if the social networks in neighborhoods are large, with many interconnected ties, the network will be more resilient and can minimize the impact of casualties that are eliminated from the network. On the other hand, a preferential (scale-free) network is tolerant to random failure but vulnerable to the elimination of a central (i.e., hub) node [26].

2. How does the severity of the disaster affect resilience of the community?
We measure severity by the number of casualties (deaths) reported at the end of each run and will study how it correlates with resilience.

3. How does the structural position of the casualties affect the resilience of the community?
It is well known in social network analysis that the relative importance of the node within the network plays an essential role. Thus the elimination of a hub or a node with high degree centrality may have the effect of breaking down network cohesion, and the elimination of a bridge node or a node with high betweenness centrality can disconnect two existing clusters. We will look at the correlation of highly central nodes being eliminated and the fragmentation of the final network.

4. How do individual beliefs affect the resilience of the community?
In this model, we explore the contribution of fear, which makes the individual move away from an infected person, and denial, whereby the individual refuses to believe the severity of the disease and behaves normally, putting himself or herself at risk.
Based on these research questions, the features that we consider in measuring community resilience are:

- the initial network topology,
- loss/casualties as represented by the number of deaths,

– the relative structural importance of casualties within the network, as measured by the average "influence" value of the dead,
– individual beliefs such as fear and denial.

4 Methods

Our ABM takes the form of semi-autonomous interdependent agents which are homophilous in characteristics (for simplicity), but individualized by their independent assignments of fear or denial of severity of a disease that, at instantiation, one among their population possesses. We maintain a constant susceptibility rate based on exposure (defined as an immediate spatial proximity to an infected other) and consistent initial spatial mapping, but the initial network structure is varied and each agent is prone to move based on its independent perception of the situation it finds itself in at each timestep. We approximate the communal tensions between individuals' desires to stay socially embedded in their communities and their desires to remain uninfected.

The product is a model in which agents move towards their socially-networked peers when they feel comfortable enough to do so, and move away from networked infected peers when they feel frightened. Agents know enough about their immediate peers to know when one is infected and decide accordingly, but may not be directly aware of other agents' status, potentially resulting in unknowing contact with an infected agent and thus inadvertent infection risks. Details of the model's specifications, instantiation conditions, and rules can be found in the attached ODD.

5 Results

5.1 Network Structure and Resilience

To investigate the role of the initial community network structure in measuring resilience, we studied the change in the number of clusters per tick-time for 100 runs comparing initial preferential and small world networks (Figure 1). The charts show that the preferential network gives rise to a much larger number of clusters at a much higher rate (i.e., more network fragmentation occurs), whereas the small world community displays more resilience by fragmenting less and enduring longer.

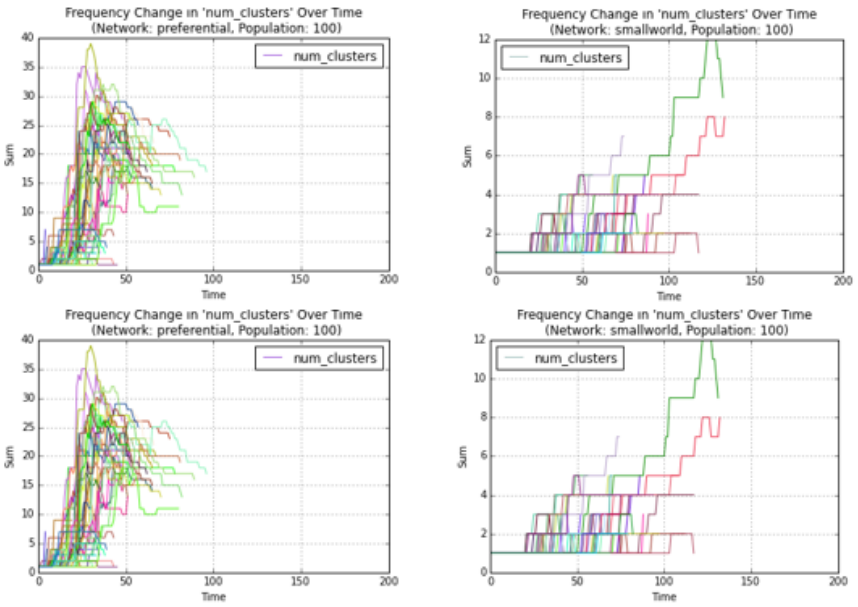

Fig. 1. Frequency change in number of clusters over time in preferential and small world net-works

Figure 2 demonstrates the known properties of our selected networks: the existence of small number of hub nodes in a preferential network produces an inequality that necessarily leaves few well-connected other when the hub dies. In our runs, we clearly see this expectation in the rate at which influential nodes die. After a few are gone from the preferential network, the average influence level quickly drops to below 0.2, but in a denser small world structure, there is more consistency closer to the 0.6 level.

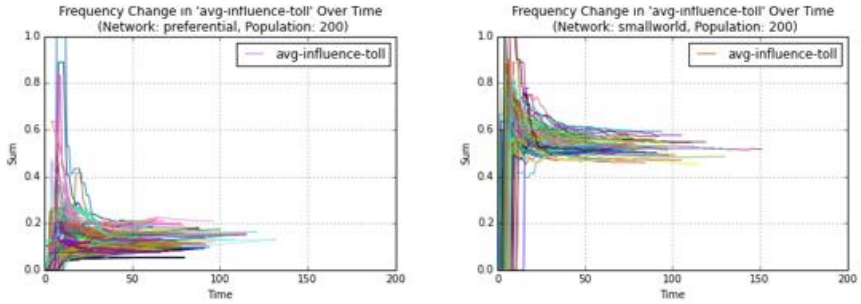

Fig. 2. Frequency change in average influence value of the dead over time in preferential and small world networks

A comparison of the change in the death toll over time for 100 runs for each network is shown in Figure 3. What emerges from these results is that the networks with higher density on the right column (small world and higher-connected random network) tend

to have slower rate of deaths, with the epidemic lasting longer. This suggests that higher density networks are more resilient, despite expectation of higher infection rates.

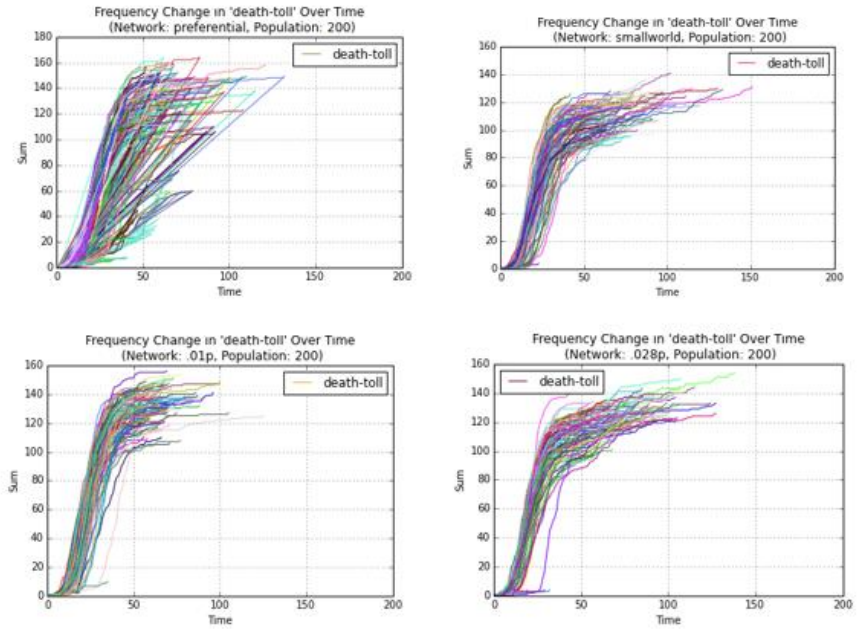

Fig. 3. Frequency change in death toll over time, initial population of 200 agents

A comparison of the number of steps taken by members of the community based on fear vs. denial are shown in Figure 4. By design, deniers who move towards the group are constrained by the communal importance of their peers, but fearful agents are restricted only by their levels of fear, so we are not surprised to this result.

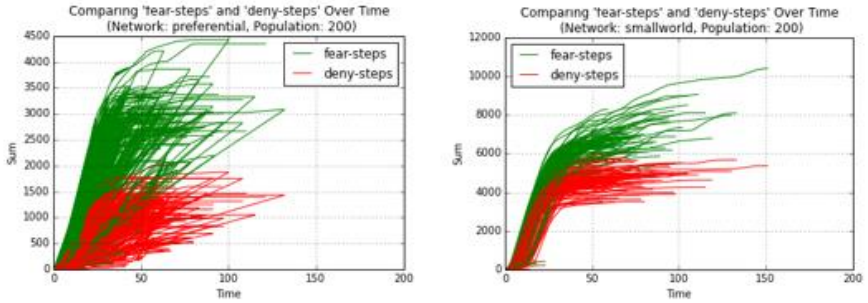

Fig. 4. Fear vs. denial over time in preferential and small world networks5.2 Severity of the Disease

Most networks initialized to one or two clusters (low link-probability random networks can have several more disconnected nodes) but as the members of the community become infected and die, the infected nodes are eliminated from the network, potentially giving rise to fragmentation of the community into distinct clusters.

A Pearson correlation between the number of casualties during a model's runtime and the number of clusters at any point in time is, on average, quite high. The 100-population 0.02 link-probability random network has a median 0.62 positive correlation. The larger 200-population equivalency with a 0.01 link-probability shows a median weak -0.144 negative correlation between these variables. The 100-population preferential-attachment network has a median correlation of 0.95 which lowers to 0.8 for the 200-population runs. Lastly, the 100-population small world network returns NaN for this metric, but the larger 200-population small world network has a correlation of 0.87. These values are represented in Figure 5.

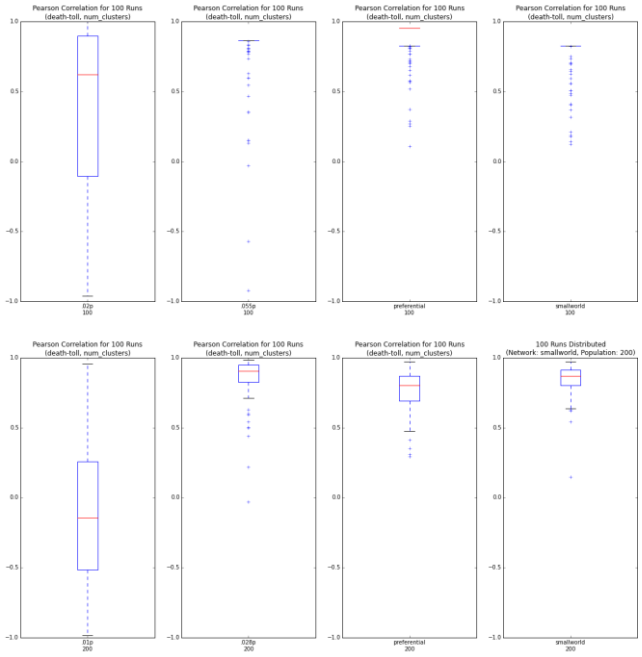

Fig. 5. Pearson Correlations for Number of Deceased and Number of Clusters

These results show that there is a strong correlation overall between the number of casualties and the amount of fragmentation in the network. We therefore expect that the more severe the disease is (i.e., propagates faster and gives rise to more fatalities), the less resilient the community's social capital will become. Where we do not see a strong correlation is with the lightly-connected (0.01 linked probability) random network. In this instance, the network starts out already fragmented with a large number of disconnected clusters; therefore a high death toll in this network has minimal impact on community resilience. In fact, the negative correlation for population of 200 reveals that mostly peripheral nodes are being eliminated.

5.3 Influence of Casualties within the Community

Pearson correlations for the average influence values of the deceased during each set of runs are depicted in Figure 6. These values are similar to those in the correlations of death-toll and number of clusters. For both population sizes, random networks with smaller probability of initial links have the largest range of observed values both inside and outside the Interquartile Range (IQR). The correlation for the smaller population and 0.02 link-probability random network is 0.67; it drops to a

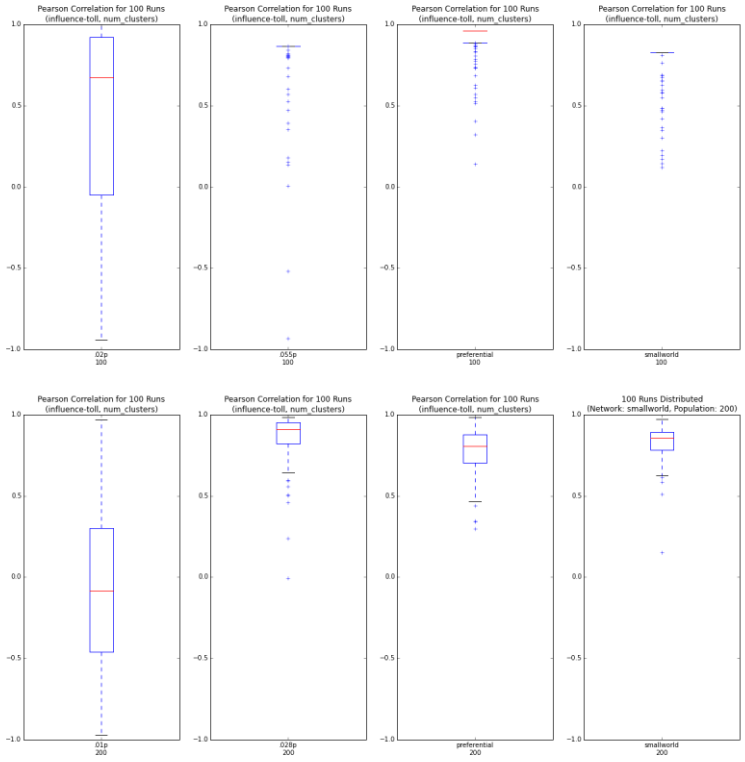

Fig. 6. Pearson Correlations for Average Influence (Deceased) and Number of Clusters

near-neutral -0.08 for the larger population model with 0.01 link-probability.

Among the 100-population runs, the 0.055 link-probability random network and the small world network do not have observed IQR in the corresponding box plots. Their average correlations are not reported. The preferential network has a 0.96 correlation. Within the 200-population runs, the 0.028 link-probability has a strong correlation with median 0.91 observed. The preferential attachment has a median correlation of 0.081 and the small world has a median correlation of 0.86.

5.4 Individual Beliefs: Denial

All the configurations report extraordinarily strong positive median Pearson correlations between deny-steps and the death-toll, as shown in Figure 7. The lowest median correlation value is 0.90 from the 100-population preferential network configuration. The highest median correlation value is 0.95 from the 200-population small world network. These results suggest that if individuals deny the severity of the disease, the death toll rises. These individuals with strong beliefs of denial are in effect taking risks, not following preventive measures, and are becoming exposed to the disease.

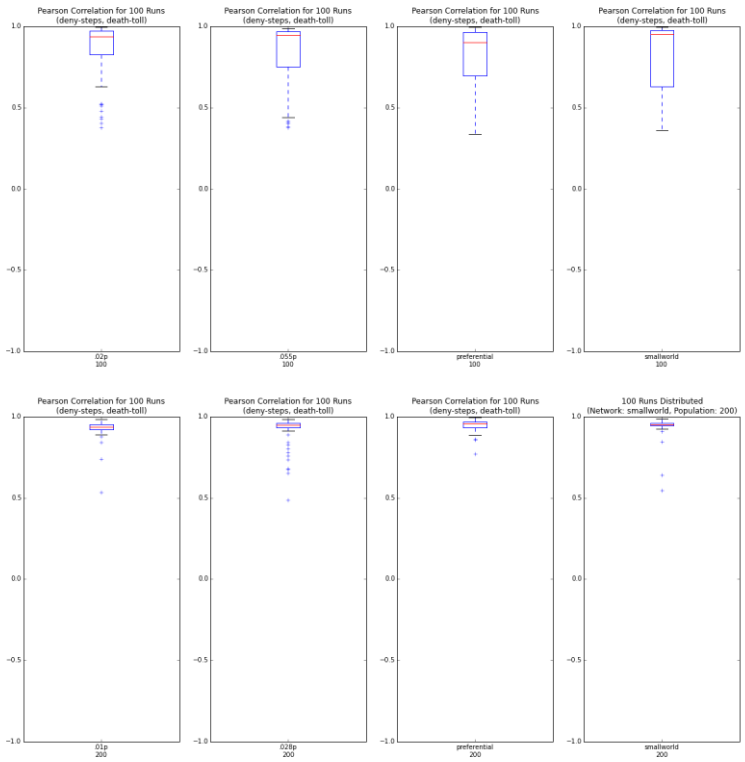

Fig. 7. Pearson Correlations for Denial-Steps and Death-Toll

However, denial correlates more weakly with network fragmentation than death toll—except perhaps in preferential networks. It may contribute indirectly by raising the number of fatalities, which in turn appears to correlate with network fragmentation. The 100-population 0.055 link-probability random network does not return a median correlation value between deny-steps and the number of clusters in the model. Other configurations range in value from 0.12 for the 200-population 0.01 link-probability

random network to 0.96 for the 200-population preferential-attachment network. Figure 8 shows the value ranges for each configuration.

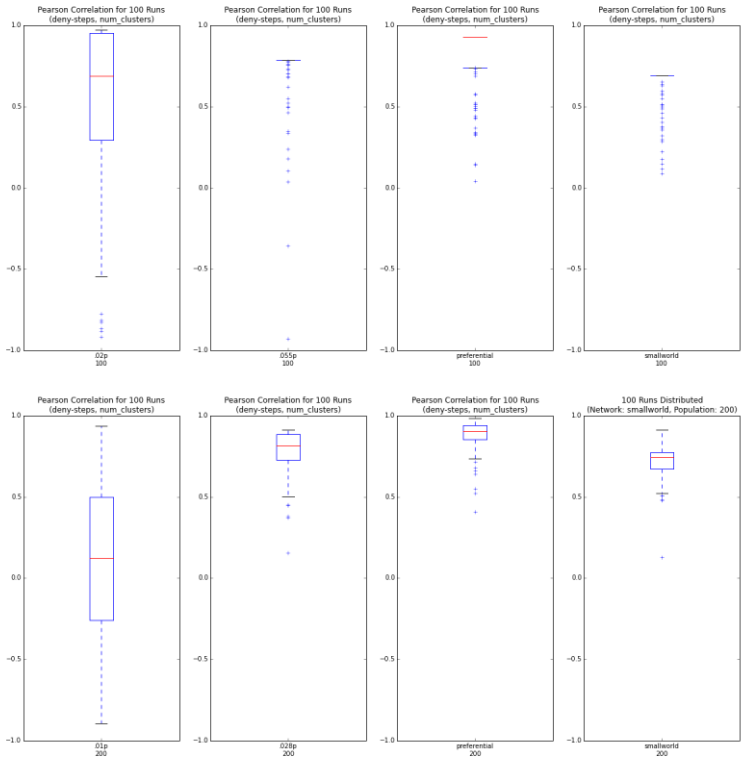

Fig. 8. Pearson Correlations for Deny-Steps and Number of Clusters

6. Conclusion

We set out to build an ABM to focus on resilience in a post-epidemic society in order to see how the original structure of the network, severity of the disaster, structural position of the casualties, and individual beliefs affect the resilience of the community. The research that went into building the model and the ABM itself proved to be very insightful in that we were able to uncover some patterns and come to interesting conclusions. In summary, we found that small world networks (and higher density networks) – in contrast with preferential networks – display more community resilience, become less fragmented, and show slower death rate. Also communities that are initially fragmented will have low resilience. In terms of the severity of the disease, to

include spread rate and number of casualties, we found that the severity level can strongly hinder community resilience. We also found that individual beliefs, such as denial and fear, can give rise to emerging behavior at the macro-level. In particular, we noticed that denial strongly correlates with a rise in death rate in the population. Finally, we found that preferential networks are more vulnerable to the effects of denial on fragmentation of the community (i.e., denial can more strongly hinder community resilience in a preferential network vs. small world network). We did not notice a strong distinction between network type in terms of the structural role and relative influence of the nodes being eliminated from the network.

The ABM explores community response to a dynamic and emerging threat, such as an epidemic, as a function of individual behaviors and connections. The results make important contributions to the modeling of social science research by combining two significant issues in agent-based modeling: (i) How individuals' cognitive biases, cultural beliefs, and strategies for making decisions shape the social environment, and (ii) How social networks evolve as an emergent outcome of social interaction and disease propagation [27].

For future work it would be valuable to consider age distribution in the community, and the role of institutions (e.g., church) in strengthening the community. We did not take into account societal institutions or the political/governmental structures. A future model could also explore the impact of available resources (such as number of medics, topology of hospitals, mitigation strategies) on community resilience.

References

1. Center for Disease Control. Ebola Virus Disease. At http://www.cdc.gov/vhf/ebola/outbreaks/2014-west-africa/case-counts.html. Retrieved on October 13, 2015.
2. Hyogo Framework for Action. 2005-2014. Building the Resilience of Nations and Communities to Disasters (HFA). At http://wwww.unisdr.org/eng/hfa/hfa.htm. Retrieved on October 13, 2015.
3. Holling, C.S.: Resilience and Stability of Ecological Systems. Annual Review of Ecology and Systematics 4. (1973) 1–23.
4. Klein, R.J., Nicholls, R.J., Thomalla, F.: Resilience to Natural Hazards: How Useful is This Concept?. Global Environmental Change Part B: Environmental Hazards, 5(1). (2003) 35-45
5. Pimm, S.L.: The Complexity and Stability of Ecosystems. Nature, 307(26). (1984) 321-326.
6. Timmerman, P.: Vulnerability, Resilience, and the Collapse of Society: A Review of Models and Possible Climatic Applications. Environmental Monograph No. 1. Institute for Environmental Studies, Toronto, Canada (1981)
7. Adger, W.N.: Social and Ecological Resilience: Are They Related? Progress in Human Geography, 24(3). (2000) 347–364.
8. Kimhi, S., Shamai, M.: Community Resilience and the Impact of Stress: Adult Response to Israel's Withdrawal from Lebanon. Journal of Community Psychology, 32(4). (2004) 439-451.
9. Walker, B., Holling, C.S., Carpenter, S.R., Kinzig, A.: Resilience, Adaptability and Transformability in Social–Ecological Systems. Ecology and Society 9(2). (2004) 5.
10. Breton, M.: Neighborhood Resiliency. Journal of Community Practice, 9(1). (2001) 21-36.
11. Adger, W.N., Huq, S., Brown, K., Conway, D., Hulme, M. Adaptation to Climate Change in the Devloping World. Progress in Development Studies 3(3). (2003) 179-195

12. Aldrich, D.P.: Building Resilience: Social Capital in Post-Disaster Recovery. University of Chicago Press (2012)
13. Adger, W.N.: Social Capital, Collective Action, and Adaptation to Climate Change. Economic Geography 79(4). (2003) 387-404
14. Aldrich, D.P.: Ties that Bond, Ties that Build: Social Capital and Governments in Post Disaster Recovery. Studies in Emergent Order, 4. (2011) 58-68
15. Rivers, C.M., Lofgren, E.T., Marathe, M., Eubank, S., Lewis, B.L.: Modeling the Impact of Interventions on an Epidemic of Ebola in Sierra Leone and Liberia. PLOS Currents Outbreaks. (2014)
16. Siettos, C., Anastassopoulou, C., Russo, L., Girgoras, C., Mylonakis, E.: Modeling the 2014 Ebola Virus Epidemic —Agent-Based Simulations, Temporal Analysis and Future Predictions for Liberia and Sierra Leone. PLOS Currents Outbreaks. (2015)
17. Piot, P., Muyembe, J.-J., Edmunds, W.J.: Ebola in West Africa: From Disease Outbreak to Humanitarian Crisis. The Lancet Infectious Diseases 14(11). (2014) 1034-1035
18. Hewlett, B.S., Amola, R.P.: Cultural Contexts of Ebola in Northern Uganda. Emerging Infectious Diseases. 9(10). (2003) 1242-1248
19. Hewlett, B., Hewlett, B.: Ebola, Culture and Politics: The Anthropology of an Emerging Disease. Cengage Learning. (2007)
20. Leone, S.: The Human Factor. Bull World Health Organ 93. (2015) 72-73
21. Onishi, N.: For a Liberian Family, Ebola Turns Loving Care into Deadly Risk. Retrieved at http://omnifeed.com/article/www.nytimes.com/2014/11/14/world/africa/in-ebola-outbreak-in-liberia-a-familys-strength-can-be-its-fatal-flaw.html. Nov 13. (2014).
22. Starecheski, L.: Life after Ebola: What it takes for a village to be resilient. Retrieved at http://www.npr.org/blogs/goatsandsoda/2015/02/20/387533191/life-after-ebola-what-it-takes-for-a-village-to-be-resilient. Feb 20. (2015)
23. Nombo, C., Niehof, A. Resilience of HIV/AIDS-Affected Households in a Village in Tanzania: Does Social Capital Help. Medische Antropologie, 20(2). (2008) 241-57
24. Frankenberger, T., Mueller, M., Spangler, T., Alexander S.: Community Resilience: Conceptual Framework and Measurement Feed the Future Learning Agenda. Rockville, MD: Westat. (2013)
25. Cutter, S.L., Barnes, L., Berry, M., Burton, C., Evans, E., Tate, E., Webb, J.: A Place-Based Model for Understanding Community Resilience to Natural Disasters. Global Environmental Change, 18(4). (2008) 598-606
26. Janssen, M.A., Bodin, Ö., Anderies, J.M., Elmqvist, T., Ernstson, H., McAllister, R.R., Olsson, P., Ryan, P.: Toward a Network Perspective of the Study of Resilience in Social-Ecological Systems. Ecology and Society, 11(1). (2006) 15
27. Bruch, E., Atwell, J.: Agent-Based Models in Empirical Social Research. Sociological Methods Research. October 23. (2013) 1-36

Overview, Design Concepts, and Details

Purpose

The purpose of this community resilience model is to explore the effectual social strain on members of a disease-exposed community starting with the existence of a patient zero at time zero. It incorporates concepts particularly relevant to social network analysts interested in structural breakdown and social behavioral scientists concerned with both the individual fear of exposure to peers who are known to be infected and the general desire to maintain social connectivity with peers. We intend this model to act as a baseline for future work and it could certainly be extended by sociologists, psychologists, behavioral health scientists, network scientists, and others.

Entities, state variables, and scales

Agents. Agents exist as representatives of humans at a 1:1 scale. These agents exist in two classes and have either one of two potential states at any given time. The first class is a basic agent with an ability to move around the stage, with its direction defined by that agent's social contacts; the other class is a medical agent with a movement guideline specifically towards an infected other agent. Both classes of agents can be in one of two states at any time and that state can conditionally change: in a healthy state, agents have no mobility restrictions (except as universally defined by model); in an infected state, agents suffer movement penalties that increase with each timestep that passes. These states have no direct impact on the direction an individual agent moves in; however, an infected agent's state may impact the movement of a social contact who intentionally tries to evade proximal contact.

All agents have three attributes, two of which remain static throughout the model. The *fear* and *denial* attributes are defined at startup as random floats between 0 and 1. They act as a proxy to the general attitude or personality of each agent and are used to determine whether an agent is willing to move towards an infected peer. Each agent also has an *influence* attribute which is determined by the number of connections that agent has relative to the largest number of connections of all agents in the model, resulting in a floated value between 0 and 1 which also changes as the network structure changes. This attribute implies that the most-connected individual has the greatest influence in the community and is determined independently of the number of connected components in the network, allowing for peer groups to exist separately from the most "influential" person in the model, but each group can have varying levels of influence which have an impact on locally-connected neighbors.

Edges. Edges between agents have no attributes. An edge simply either exists between two agents or it doesn't. Edges are destroyed when one of the agents responsible for its existence dies. Edges are not created except by medical-class agents who do not otherwise have an infected neighbor (effectively ensuring that every medical agent in the model is trying to mitigate the effects of the disease throughout the model's runtime).

Spatial units. This model has a NetLogo default size of 32x32 and maintains an agnostic but bounded environment: agents can roam in any direction (effectively grid-less) and all parts of the world are identical in function (none), but the world does not wrap-around, meaning that the perimeters are like impassable walls to the agents.

Process overview and scheduling

There are two primary activities performed at each timestep (after initialization): first, structural calculations and agent status updates are performed; second, agents make decisions about where and how quickly to move.

Checks and updates. At each timestep, several checks are made before any agents move. First, network centralities are updated to reflect any structural changes during the previous timestep; the model tries to measure averages for the following metrics, calculated by NetLogo's networks extension: betweenness, eigenvector, closeness, and clustering coefficient. Second, medical agents are activated and a check is made to determine whether the medical agent has a connection with any infected agent, creating a new connection if none exist. Thirdly, agents who were infected in the previous timestep are activated and several checks are performed: a) whether the agent is still infected (not changed during last timestep), b) whether the infected agent infects another agent within close spatial proximity, c) whether the agent dies of the infection, and d) whether the agent is near a medical agent and is suddenly cured by that medical agent.

Movement. After the above checks and status changes are performed, each agent is activated in some order determined by NetLogo (in a preserved order by unique ID determined at instantiation) so that the agent may move. Movement is impacted by the attributes defined above (fear, denial, influence of peers) as well as infected status: infected agents are quickly handicapped and immobilized according to the rules described in *infected movement* below, but otherwise attempt to move according to the same rules as healthy agents. Agents are given individuality insofar as the selected direction is impacted by the agent's fear versus denial values and direction is impacted by its social peers; the rules governing direction and movement speed are set universally, however.

When activated, a healthy non-medical agent assesses its connected neighbors (connected by social network, not spatial proximity) to determine if any of them is infected. If none are infected, the agent will move according to the *safe movement* rules defined below. If, however, one or more connected neighbors are infected, the *safe movement* rules apply if the agent has a denial value greater than its fear value; otherwise, the *fear movement* rules are applied.

Medical agents generally move according to the rules described by *medical movement* below. They may also be infected and, if so, will suffer movement penalties in the same way other agents do and will also lose their ability to cure other infected agents.

Lastly, as community structure breaks down, agents may become isolated (no connections). These agents also move according to the *fear movement* rate, except they move in a random direction, not influenced by anything defined by the model.

Safe movement. The term "safe" is more aligned to the agent's perception of its movement decision than to the calculated reality of the outcome. This type of movement sets the agent's direction as the averaged x- and y-coordinates of all its neighbors (regardless of infected status) at a rate equal to (max-movement speed * average influence of connected neighbors).

Fear movement. An agent that is connected to at least one infected other and has a fear value greater than a denial value will perform this movement. The direction it moves is exactly opposite the averaged x- and y-coordinates of all connected neighbors who are infected, in an attempt to maximize its distance (and thus reduce perceived risk of infection) from infected neighbors. The movement rate is equal to (max-movement speed * fear).

Medical movement. A medical agent moves at a typical rate defined by the *fear movement* above (max-movement speed * fear) but moves towards the nearest infected agent with which it shares a direct connection. Unlike other agents, there is a conditional modification to this movement: if the nearest infected agent is at a distance less than the calculated rate, then that is the distance moved, so that the actual movement is the lesser of the two: (max-movement speed * fear) or (distance to nearest infected and connected agent). Thus, even if the infected agent moves, the medical agent will remain close enough to potentially cure it. If the medical agent has no connections that are infected (due to a programming logic error in which infected agents can die after medical agents establish their patient list), it performs a *safe movement* instead.

Infected movement. Agents which are infected will follow all the same rules as agents that are not infected, but suffer a movement penalty so that max-movement speed is reduced by the number of timesteps spent infected. Thus they move at a rate equal to ((max-movement speed – time infected) * modifier).

Isolated movement. Isolated agents choose a random direction and move at a rate equal to ((max-movement speed – time infected) * fear), adding noise to the model, as though the agent goes about its business without up-to-date information about the community.

Design concepts

Emergence. Contrary to many network studies, we're interested in the decomposition of network structures, the opposite of the emergence of network structure.

Adaptation. Each agent has a *fear* and a *deny* attribute. These are meant to represent an interaction between belief and personality characteristics, resulting in sometimes isolating, sometimes stubborn and risky behaviors. Each attribute is independently a value

randomly selected between 0 and 1. When an agent has a neighbor who is infected, it evaluates its own fear and deny values to decide how it will move about the world.

When agents move, they have three movement types: *random-move*, *safe-move*, and *fear-move* (see above on Process Overview and Scheduling). Although the attributes are static, the network structure is not. Agents assess their social environment and make decisions accordingly at each timestep, altering the geospatial locations of agents in the model and affecting levels of risk of exposure to infection.

Figure 9 shows an example initial network and final outcome. The circular positioning is initialized with the model (see the next section on Initialization for more information). final network after all infected members of the community have either died or been cured. There still remains one large connected network with two smaller connected networks. The other agents have been disconnected from the main community due to the casualties. However, since the movement in the model is spatial, a number of clusters have been formed (without overt social ties) as the members have moved towards each other.

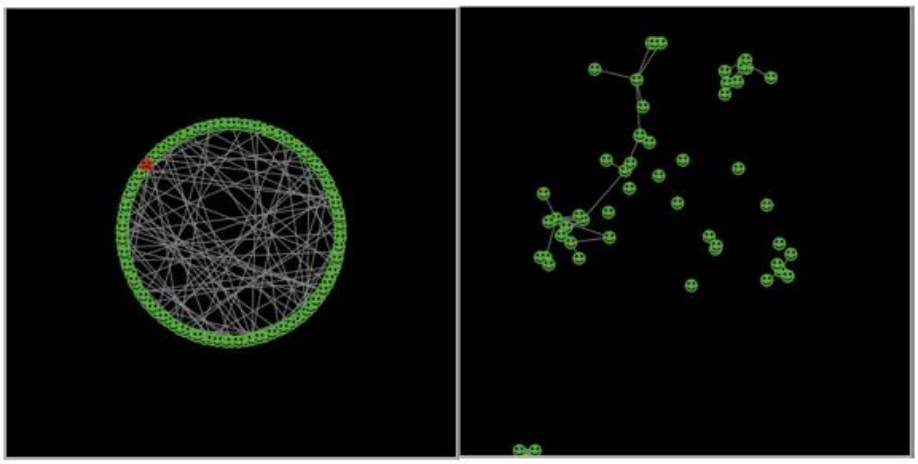

Fig. 9. Agent-Based Model snapshots with initial preferential network, population = 100 (left) and resulting network at the end of system run (right)

Interaction. Interaction between agents isn't direct. Agents that are classified from the beginning as *medical* are obligated to try assisting the infected agents and, as a result, will move directly towards its closest infected peer. By being in proximity to an infected individual, we assume interaction and assign a probability of curing the infected. Similarly, an infected individual has a probability of infected nearby others, but does not interact with any of them directly.

Collectives. Agents are initialized in generally one—sometimes two—connected networks. As the individuals die and are removed from the network, the network structure(s) is altered. The entire population's most-connected agent is always observed as relative to the entire population, while each agent queries the relative proportion of ties

each of its own neighbors (as well as the infected-status of its own neighbors) when deciding how to move in the environment. In this way, agents exist in, influence, and are influenced by local connectivity, but do not escape or have perfect knowledge of the whole world.

Observations. As the model runs, each agent reports its own calculated centrality value for each of three centrality metrics: *betweenness*, *closeness*, and *eigenvector*. Unfortunately, a limitation of NetLogo's networks extension is that the betweenness metric is not normalized (reduced to a proportional value between 0 and 1). Another is that eigenvector cannot be calculated for each cluster on the stage; if the existing agents do not collectively comprise a single cluster, attempts to measure eigenvector result in a "N/A" value. Each agent also knows the number of edges it possesses (*degree* metric) and compares this value against the largest number of edges any existing agent possesses at each time-step, resulting in the possession of edges equating to a relative proportion of the maximum in existence. This value is converted to an *influence* value (between 0 and 1) for each agent which changes as the network structure changes.

We calculate and store the global averages for each of the above measures at every time-step. Additionally, we monitor the number of *clusters* (independent groups of connected agents) over time. The model typically starts with a single cluster (occasionally two, in the case of randomly-generated networks) and fragments as agents become infected and die. Along with the above variables, each agent has a *clustering coefficient*, which represents how well-connected its neighbors are to one another. The global average of this variable is captured at every time-step, as well.

Initialization

The model is initialized with a pre-selected network type (*random*, *preferential attachment*, or *small world*) and predefined *population* set to either 100 or 200 agents in a circle-layout with radius size 8. Exactly 1% of the population (or minimum of one agent) becoming infected and up to 1% of the population classified as "medical." Exactly one or two agents are infected before the first round begins for every run. However, it is not guaranteed that a medical agent will exist in each run. It is possible for the infected agent-zero to be a medical agent. Agents cannot become medical after the model is initialized.

Although the model starts with just one or two infected agents, infection can spread to nearby agents. This captures the risk of near-proximity exposure for uninfected individuals and interacts with individuals' personal beliefs about the risks of exposure (discussed below). The probability that an infection will spread to another agent is a static rate determined by a slider on the user interface. Our simulations maintain a probability of 10%.

The social network connections of our model are assigned at setup and do not change except under two conditions: 1) if an agent dies, its links (if any) to other agents disappear, changing the overall network structure; 2) as discussed below, there is a single mechanism through which medical-type agents create links with agents who are infected, resulting in a slight modification of the network structure and therefore how decisions are made by agents.

Input data

No external data sources were used for this model.

Submodels

There are no submodels to include here. Formulas used in the model are described in the section titled “Processes overview and scheduling.”