SAND 2011-3189 C

Extending Opinion Dynamics to Model Public Health Problems and Analyze Public Policy Interventions

Thomas W. Moore, Patrick D. Finley, John M. Linebarger, Alexander V. Outkin, Stephen J. Verzi, Nancy S. Brodsky, Daniel C. Cannon, Aldo A. Zagonel, and Robert J. Glass
Sandia National Laboratories
Complex Adaptive Systems of Systems (CASoS) Engineering Initiative
tmoore@sandia.gov

The public health community is recognizing the importance of social network dynamics in analyzing diseases correlated with behaviors such as tobacco and alcohol use, substance abuse, and poor nutrition and inadequate physical activity. These behaviors are driven in part by opinions that individuals hold regarding products, behaviors, and lifestyles. The opinions and behaviors of individuals are influenced by their personal social networks as well as exogenous components such as advertisements. We extend the basic opinion dynamics model to include two processes important for analysis of diseases caused by unhealthy behaviors. The first is an antagonistic reaction that drives individuals further apart in opinion space; the second is the addition of hysteresis representing the constraint addiction places on an individual's behaviors. We apply this extended model to consider tobacco use within a community and various approaches to influence its prevalence, including advertisements and health-related educational campaigns. We examine the roles of advertising strength, the strategic importance of tolerance, and how hysteresis in the behavioral function influences tobacco usage within a community. Finally, we show how spatially and temporally local results can act as inputs to a population-wide, long-term system dynamics model. This allows for the examination of the impact of interventions on future mortality.

1 Introduction

Many chronic diseases can develop or progress due to behavioral choices of individuals. These diseases, including some types of heart disease, cancers, and many metabolic disorders such as diabetes often correlate with behavioral components such as diet and physical activity, smoking, and alcohol and substance abuse [Hjermann et al., 1981; Single et al., 2000; Stampfer et al., 2000]. Studies have demonstrated social-network based clustering effects for these behaviors that are similar to those shown for communicable diseases [Christakis & Fowler, 2007, 2008; Rosenquist et al., 2010]. Analysis of social network mediated interactions has proven fundamental to the understanding of contagious disease epidemics, such as influenza [L. M. Glass

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

& R. J. Glass, 2008]. Although chronic diseases themselves are often not considered communicable diseases, similar propagation of behaviors through social networks may be a causal factor in patterns of these diseases in the population [Smith & Christakis, 2008].

As individuals interact with others in their social networks, they exchange beliefs, ideas, and opinions in both direct and indirect ways. As an example, simple discussion of ideas and opinions between members often leads to some individuals convincing others to modify their opinions about the concepts or beliefs discussed, while extended social interaction can result in each individual gradually modifying his or her opinion toward a more consensual view in search of “common ground.” This phenomenon may be seen as an application of what social psychologists have identified as structural balance theory, which states that a positive affective relationship between two individuals will tend to lead them towards similarity in their affective relationships to a third individual or concept [Dorwin Cartwright, 1956]. In addition to these relationship-mediated means of opinion and belief exchange among individuals in social networks, media sources can influence the opinions held by members of a community through elements exogenous to the immediate social networks via mechanisms such as television, billboards, and radio broadcasts. These exogenous elements may act directly, as in the case of cigarette advertising, or indirectly, in the form of behavioral modeling and influencing perceptions of social norms.

Direct and indirect exchange of opinions and ideas within social networks may result in changes in individuals’ behaviors. To the extent that an individual’s actions are influenced by their opinions, it can be seen that changes in opinions may result in changes in behaviors. If opinions can be seen as propagating through networks, and opinions influence behaviors, then one of the most direct observable results would be a tendency of the resulting behaviors to cluster in social networks, forming smaller sub-networks of individuals with similar opinions and behaviors.

Opinion dynamics modeling is a recently developed family of approaches for the analysis of social influences on individual opinions and the emergence of resulting community-scale patterns. These models have been developed by the statistical physics community and are grounded in Ising models of particle spin alignment in lattices [Castellano et al., 2009]. Such models encompass significant variation in approaches: binary, discrete, or continuous opinion values; unstructured, linear, lattice, or complex network topologies; and random or averaging interactions. However, all opinion dynamic formulations share common theoretical roots, and all generate clusters of individuals sharing similar opinions based on local rules governing individual interactions. In these models, a set of individuals are used to populate a community, and are seeded with initial opinion values. Each individual updates her opinion based on interactions with her neighbor(s). These interactions are potentially governed by network topologies, randomness, and similarity in individual opinions.

2 Opinion Dynamics Model

Our model extends a widely used model introduced by Deffuant et al. [Deffuant et al., 2000]. This approach, frequently referred to as the Deffuant-Weisbuch (DW) model, was initially constructed using randomized interactions among individuals in a well-mixed population. Individuals are assigned a random opinion, taken as a value on the continuous interval [0, 1] drawn from a uniform distribution, and a tolerance threshold, ε. In this model, the tolerance threshold of an individual limits the number of interactions that will result in an opinion change. This value can be thought of as a measure of uncertainty or open-mindedness about a given issue, in which an individual is willing to "listen" (that is, marginally update her opinion based on the opinion of her neighbor). If the difference between her opinion and that of her neighbor exceeds her tolerance threshold, she will be unwilling to listen to her neighbor on the issue, and no change in her opinion value will occur.

We apply the DW model to directed social networks. Directed social networks can represent types of relationships often characterized as nominations, for example, as gathered in a survey asking individuals to name their closest friends. Although the original work and some later extensions concentrated on reciprocal exchanges of opinion, friendship networks are often represented using directed relationships[Scott, 2000]. In addition, some of the empirical studies considering network-based properties of tobacco use have identified correlations incorporating directionality[Christakis & Fowler, 2008].

We interpret the dynamics of the model to represent overall social influences from all nominated individuals rather than discrete pair-wise interactions. That is, the model considers the continuous interactions between friends, rather than the discrete exchanges that would occur in a deliberation on a particular subject. In the case of a node with multiple out-edges (for instance, an individual who has named more than one person as a friend), we average over the opinions of the connected nodes. Our equation becomes:

This averaging effect is similar to the one proposed in some alternative implementations of the opinion dynamics model [Hegselmann & Krause, 2002].

Here, is the cardinality of , and is the set of all neighbors of whose connecting edge points from and whose opinions fall within the tolerance threshold, determined by evaluating the absolute value of the difference in opinions against the tolerance value:

The model can be viewed as a social network of individuals seeking to gain consensus with their neighbors. At each time step, each node of the social network graph adjusts its opinion value to a value closer to the mean of its neighboring nodes. When this process is applied across all nodes of a network, opinions of nodes in certain portions of the graph will tend coalesce to common mean values, with the number and average size of the clusters primarily determined by the constraining tolerance variable ε [Weisbuch et al., 2002]. Tolerance constrains interactions and encourages isolation and cluster formation by setting an upper bound on the number of interactions that result in a change of opinion. The portions of the graph whose nodes display similar opinion values define opinion clusters. As shown in Figure 1, over repeated time steps the opinion dynamics model causes the social network to shift from isolated nodes with randomly distributed opinions to clusters of neighboring nodes sharing a common opinion. With a high tolerance value (shown in panel B), opinions converge to a single consensus value. Lower tolerance values cause heterogeneous clusters of opinion to form (shown in panel C).

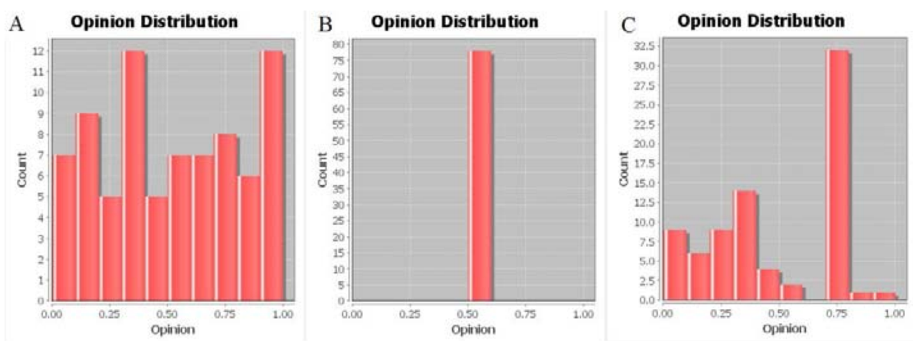

Figure 1: Histograms of opinion distributions for a 75-node scale-free network showing (A) initial opinion distribution, (B) final steady-state distribution with tolerance=0.5 and (C) with tolerance=0.2

2 Network Topologies

Social network topologies control which nodes are direct neighbors to a given node. While survey-based social networks are useful for determining social relationships within a community, such surveys must be well constructed and the resulting responses carefully analyzed. Random networks constructed to resemble those obtained from surveys are a useful alternative that allows many different social network structures to be investigated efficiently.

Random networks form the basis for our simulations using opinion dynamics to model public health issues. Scale-free networks are often created using the method of preferential attachment. Preferential attachment network construction generates topologies that exhibit a power law distribution of node degree [Barabási & Albert, 1999]. Scale-free topology has been repeatedly discovered in a wide variety of phenomena, including computer networks and websites, protein interactions in cellular physiology, and social networks representing friendships, advice-seeking, and sexual relations [Albert & Barabási, 2002]. Preferential attachment explains some, but not all, topology in friendship networks [Jackson, 2008]. We modify a scale-free network to include a proportion of edges between randomly selected nodes, resulting in a network that is predominantly constructed using the Barabasi-Albert model of scale-free network construction, with a smaller proportion of edges determined by an Erdos-Renyi random process.

3 Antagonism

A variety of interactions can be modeled with DW opinion dynamics models. The original definition of DW opinion dynamics specifies two types of potential interactions between individuals: positive interactions, in which the individuals' opinions move closer to one another, and neutral interactions, in which the opinions are considered too far apart so that no adjustment takes place. Although these two possibilities capture a wide range of potential interactions, some researchers have recently added a third possibility: a negative interaction that drives the opinions of the individuals further apart. The potential for this antagonistic response is attributed to the "ego-involvement" of the individual agents according to an interpretation of Social Judgment Theory [Jager & Amblard, 2005].

If an interaction occurs between individuals whose opinions differ by an amount greater than the antagonism threshold value, the resulting change in opinions is identical in magnitude to the original equation, but opposite in sign:

Here, is the set of all out-degree neighbors of whose opinions fall within the bounds of antagonism.

Addition of antagonistic responses to simple DW dynamics enables us to simulate a more complete range of opinion-dynamics interactions between agents:

1. Consensus with entities adopting new opinions closer to their neighbors whose opinions are already similar
2. Indifference with entities not affected by opinions of their neighbors where opinion differences exceed the tolerance threshold epsilon
3. Polarization with entities adopting widely divergent opinions when opinion difference is greater than the antagonism threshold.

4 Analyzing Smoking Using Opinion Dynamics

We consider the case of cigarette smoking in a community as an illustration of an application of these ideas to public policy analysis. In addition to smoking being the leading cause of preventable deaths in the United States, responsible for 18.1% of total deaths in 2000 [Mokdad et al., 2004], multiple researchers have demonstrated strong correlations between smoking and social network relationships [Christakis & Fowler, 2008; Galea et al., 2004; Valente, 2003].

Youth experimentation with smoking is primarily catalyzed by psychosocial motivations, especially aspirational components including rebellion and an assertion of independence and adulthood [Jarvis, 2004]. The tobacco industry capitalizes on these aspirational components by targeting brands to specific socio-economic segments and designing marketing campaigns that create associations between these aspirational components and tobacco products [Jarvis, 2004; Ling & Glantz, 2002; Pierce et al., 1998].

We interpret the opinion value of an individual to represent that individual's opinion about smoking in this opinion dynamics investigation of tobacco use. Ideas such as "Smoking helps people control their weight" and "Smoking is cool" could contribute to a favorable opinion about smoking. Alternatively, ideas including "Smoking causes lung cancer" and "Second-hand smoke is dangerous" could contribute to a non-favorable opinion about smoking. We interpret the opinion value for a given agent to be an aggregate value representing the agent's belief in all such ideas. Using a continuous range of opinion over [0, 1], we interpret an opinion value of 0 to be extremely anti-smoking, an opinion of 1 to be very favorably disposed toward smoking, and opinions in the range [0.45, 0.55] to be essentially neutral on the topic.

5 Opinion-Behavior Mapping

Opinions are of interest in this investigation because they are assumed to affect behavior. For purposes of simplicity, our model proposes a simple step function with the value of the behavior being either *true* or *false*. We set an initiation threshold; when an agent's opinion exceeds the threshold value, the agent initiates the behavior. The initiation threshold can be interpreted as a subjectively assigned measure of utility of smoking to the individual. If the perceived utility cost of a behavior is high, the individual needs a higher opinion about the behavior than they would if the cost was relatively low. Cost is interpreted as not only monetary cost but also convenience. In this model, cost does not include perception of harm or social costs. Negative concepts associated with smoking contribute instead by lowering the opinion value. Initiation thresholds for smoking could be raised by increasing the purchase price for a pack of cigarettes, but also through indoor smoking restrictions or age-based point-of-sale restrictions, either of which make acquiring or smoking cigarettes more difficult.

We apply hysteresis in the function that maps opinion to behavior when the behavior of interest has a physiologically or psychologically addictive component. This formalizes the notion that addiction compels an individual to maintain the behavior even when her opinion falls below that which would cause initiation. This lack of correlation between falling opinion and smoking cessation can be seen in surveys indicating that 70% of smokers stating that they wish to quit, with 33% attempting to quit each year with a less than 10% success rate unless additional assistance in quitting is received [Rigotti, 2002]. We allow for various degrees of addictiveness of products and addiction of individuals by setting a cessation threshold to some value less than the initiation threshold. In the case of cigarette smoking, the cessation threshold can be increased, thus lowering the effects of addiction, through the use of support groups and nicotine replacement therapy. Figure 2 illustrates this model, showing non-smoking behavior occurring until the opinion passes the initiation threshold in the direction of increasing opinion. Once smoking initiates, it will continue to occur even with decreasing opinion until the falling opinion crosses the cessation threshold.

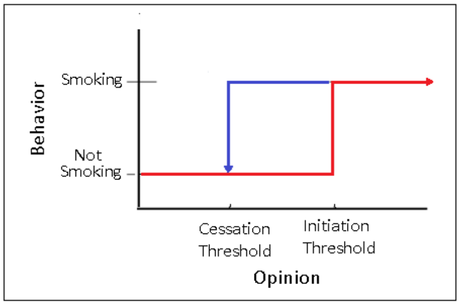

Figure 2: Graph illustrating addiction hysteresis. Red path shows that no smoking occurs until initiation threshold is reached. Blue path shows that smoking behavior continues until cessation threshold is reached.

6 Scenarios

We apply our model to the analysis of three different scenarios to illustrate the utility of modeling for understanding the impacts product advertising and public health policy measures. First we examine how advertising and educational campaigns influence opinions and resulting behaviors in social networks. Next we examine information campaigns which modify individual tolerances at strategic locations in the network. Lastly we examine the effects of actions which could shift threshold values for initiation and cessation.

6.1 Using Advertising and Education to Influence Opinion

We model the effects of this information flow by allowing the model to determine population clusters resulting from different levels of initial opinions and tolerance thresholds. Agents external to the social network representing, for example, industry and health advocacy groups, attempt to modify the behavior of individuals through the use of advertising and educational campaigns. We model information flows from these external sources as media nodes that inject new opinion values to selected individuals within the network. We have adopted a terminology convention to differentiate efforts by the tobacco industry from those of public health groups. We refer to industry efforts to promote the smoking as “Advertising.” Conversely, public health messaging campaigns to counter the behavior or encourage healthier alternatives are denoted as “Education” or “Countermarketing.”

We consider an advertising or educational campaign to be a specialized media node having only in-edges. This indicates that information flows from the specialized media node to other connected nodes. That is, an educational poster at a bus station may influence people, but they cannot, in turn, directly influence the opinion of that poster.

Advertising, education, and countermarketing campaigns seek to communicate with the most influential members of a social network in hopes that the message will then propagate from the influential individuals to others. An advertising or educational node can be configured to connect with influential nodes in the social network by targeting nodes with specific network properties. For example, social network members who regularly communicate with many others in the network are represented by network nodes having greater in-degree (number of in-edges). By targeting these individuals whose network node importance measures are high, the model encapsulates accepted marketing and public relations concepts. To effect behavioral change across the social network, our model allows us to examine the effectiveness of different network-node importance metrics such as in-degree centrality (proportional to the number of in-edges) and betweenness centrality (proportional to the number of paths through the network on which the node lies). Advertising and educational nodes attempt to influence the network as a whole by injecting an opinion value into these important connected nodes using the same modified DW opinion dynamics mechanism introduced above.

An advertising campaign can attempt to raise opinions about smoking through positive associations. We can model such a campaign as an attempt to influence the network strongly by projecting an opinion value close to 1.0 to important nodes. Similarly, an educational campaign can attempt to dramatically lower opinions by espousing an opinion value close to 0.0. These extreme values, however, can fall outside the range of tolerance for the individual nodes to which they are attached, either failing to influence or, as a result of antagonism, pushing the individual and the network in the opposite direction.

We find that the ability for a node to influence the network via opinion propagation is primarily determined by an individual's PageRank, a centrality ranking algorithm closely related to Eigenvector ranking. PageRank emphasizes importance as determined by random walks through the graph. Using the PageRank method, the importance of a node is determined not only by the number of nodes pointing to it, but also the relative importance of those nodes [Brin & Page, 1998]. An example of the potential effectiveness of media nodes is illustrated below (Figure 3). Using a scale-free network with 17 nodes (N=17), a tolerance threshold of 0.3, and initial opinions seeded randomly from a uniform distribution on [0, 1], the network converges to a steady state consensus opinion of 0.55. Attaching a single media node with an opinion of 0.35 to the original network changes the steady state consensus opinion value to 0.37.

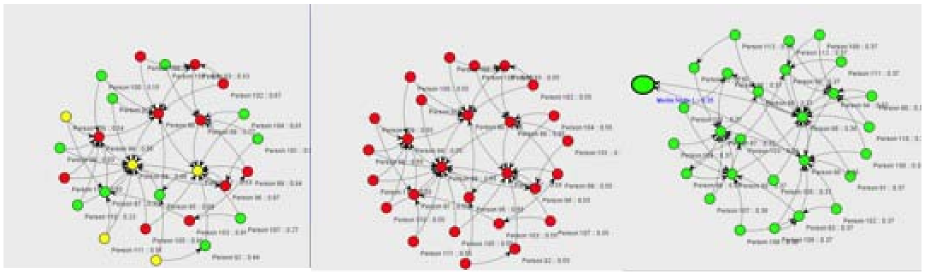

Figure 3: Left Panel: Initial social network. Center panel: Same social network after running opinion dynamics algorithm. Right Panel: Same network after attaching education node and then running opinion dynamics algorithm. Red nodes have a positive opinion towards smoking (opinion ≥ 0.55), green nodes have a negative opinion towards smoking (opinion ≤ 0.45), and yellow nodes have a neutral opinion ($0.45 \leq$opinion ≤ 0.55).

The ability of a single campaign to influence the network is constrained by the tolerance values of the individuals in the network. A network composed of individuals with a low tolerance to opinions of their neighbors' results in many opinion clusters with few individuals in each cluster. Social networks composed of individuals with high tolerance thresholds typically coalesce to one to two dominant opinion clusters. This implies that advertising campaigns conducted on a network with low tolerance can change opinions in isolated opinion clusters, but a higher tolerance network is required for injected opinion to propagate throughout the network. Because this model uses tolerance to represent general open-mindedness or uncertainty regarding an issue, tolerance is not considered to be biased in either direction. An individual's tolerance extends equally in both the pro-smoking and anti-smoking directions, meaning that the analysis of the impact of tolerance on advertising campaigns applies equally to the impact of tolerance on educational and countermarketing campaigns.

A campaign attempting to generate network consensus can force a large part of the network to a moderate opinion value (e.g., [0.40, 0.60]), but has more difficulty bringing a shift to a more extreme consensus. The opinion promulgated by advertising or educational campaigns must be within the tolerance threshold of the targeted individual for the message to effectively shift opinion. Messages that are outside of an individual's tolerance are interpreted to be extreme and are ignored or serve to drive that individual's opinion in a direction opposite of that intended by the campaign. A moderate message is more likely to be within the tolerance threshold of more members of the social network, and thus can be quite effective in shifting opinion to central values. It's also possible for moderate-valued campaigns to decrease extreme opinions from their own side; a campaign promoting a moderately favorable opinion about smoking with an opinion value of 0.60 may end up dragging down network clusters that would otherwise converge to a higher opinion value. The converse possibility holds for campaigns promoting a non-favorable opinion. Thus to shift opinion to extreme values, the injected message must be more extreme than the desired final opinion value. Additionally, extreme messages are more likely to be outside the tolerance of individuals holding moderate or opposite opinions, thus fewer nodes are available to be influenced with a truly extreme campaign.

Opinion toward a behavior can change considerably without affecting outwardly-directed behavior. The opinion of an individual must exceed the initiation threshold for the characteristic behavior to begin and must fall below the cessation threshold for the behavior to cease. Unless the campaigns shift individuals' opinions across one of these thresholds, behavior will not be initiated or stopped. A campaign can therefore be effective in bringing about an opinion shift, but be ineffective in bringing about a significant change in the behavioral regime. In our model, this means that an individual may become favorably inclined toward smoking but be unwilling to bear the financial and convenience costs to adopt the behavior, or conversely they might develop a non-favorable opinion about smoking, but not sufficiently so to overcome their addiction.

Compound advertising or educational campaigns consisting of multiple messages working in concert can increase effectiveness over that of either message alone. A compound intervention employs multiple campaigns. The initial campaign pushes a moderate opinion that is within the tolerance of individuals holding anti-tobacco opinion. This initial campaign serves to shift the opinion of these anti-tobacco individuals to a more moderate position. The follow-on campaign then applies a more strongly pro-tobacco message, which can then move the already biased network to the desired value. This complementary effect can be used to generate a widely held consensus at a value well above or below the initial mean opinion. Figure 4 illustrates the effects of complementary ads, showing the mean results of 204,000 runs on randomly generated social networks.

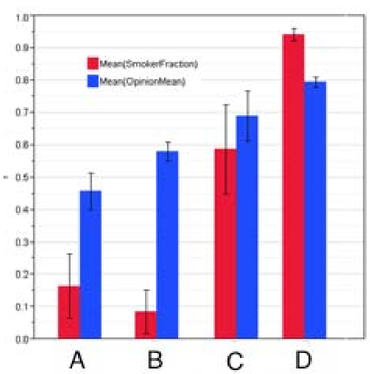

Figure 4: Results of complementary advertising analysis. Red bars indicate fraction of population who smoke, using an initiation threshold of 0.65 and a cessation threshold of 0.35. Blue bars indicate mean opinion on tobacco for population (0.0 = Unfavorable, 1.0 = Favorable). Error bars indicate 95% confidence interval of the mean. Results shown for no advertising (A), mild ad (B), strong ad (C) and strong/weak complementary ads (D).

We analyzed the ability for advertising nodes to influence a scale-free network of 250 nodes (Figure 4). Advertising nodes were connected to the top ten most important nodes in the network determined using the PageRank method. The mild advertisement opinion was set to 0.65, and the strong advertisement was set to 0.85. The control results indicate network behavior in the absence of advertisements. A mild ad acting alone was able to raise the average opinion of the network from approximately 0.46 to approximately 0.58. However, a side effect of the mild ad is to decrease the opinions of individuals below the initiation threshold. This results in the unintended side effect of decreasing the smoking fraction from approximately 16% to approximately 8%.

A strong ad acting alone was able to raise the average opinion significantly higher, albeit with an increase in variability as shown by the error bars in Figure 5. The average opinion increased from the baseline of approximately 0.46 to approximately 0.69, while the smoker fraction increased from approximately 16% to approximately 58%. The strong ad was thus significantly more effective than the mild ad, both in changing opinion and in changing behavior.

The strongest observed effect comes from combining the two advertising strategies. With both the mild and the strong ads connected, mean opinion was raised to approximately 0.79, while the average smoker fraction increased to approximately 94%. Complementary advertising campaigns can run consecutively, with the mild campaign preceding the strong one, or concurrently. In concurrent campaigns, individuals who are initially unaffected by the strong campaign can have their opinions modified by the mild campaign, and eventually move close enough to the strong campaign's position that they become affected by it.

6.2 Using Advertising, Countermarketing, and Education to Affect Tolerance

Tolerance in opinion dynamics indicates the receptivity of an individual to a differing opinion. Tolerance is sometimes termed "lack of certainty" about one's own opinion. Low tolerance values thus effectively limit the breadth of opinions individuals are willing to incorporate into their own. An advertising or educational campaign can affect tolerance if it is designed to adjust an individual's willingness to listen rather than affecting their opinion relative to a product or behavior directly. For example, a claim that expert scientific opinion remains divided on a subject (for example, the effects of secondhand tobacco smoke) could lead to an increased tolerance among some individuals, producing a willingness to give more credence to opposing opinions. A tolerance-based campaign might conversely bring about a change in opinion by raising questions about bias and deception by tobacco companies in the presentation of evidence, as was done by the "truth" campaign [Farrelly et al., 2005].

Our research indicates that the ability of an advertising or educational campaign to affect the network is grounded in the tolerance values of the nodes with the highest betweenness centrality. The betweenness of a node is proportional to the number of shortest paths on which it lies, with a greater number of shortest paths running through a node contributing to a higher betweenness rank.

We analyzed the ability of an advertising node to influence the opinions and behaviors in a scale-free network of 250 people (Figure 5, below). The network was initialized with a uniform distribution of opinions on the range [0, 1] and the advertising node was propagating an opinion of 1.0 (Very favorable to tobacco). Baseline tolerance for nodes was set to 0.50. The advertising node was connected to the four nodes with the highest PageRank values. Tolerance values were varied for the six nodes with the highest betweenness rankings.

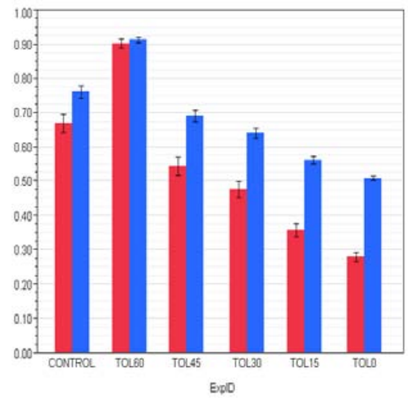

Figure 5: Results of tolerance-based advertising analysis. Bars represent mean final values for opinion dynamics model run over 204,000 different generalized social networks of 250 nodes. Red bars indicate fraction of population who smoke. Blue bars indicate mean opinion on tobacco for population (0.0 = Unfavorable, 1.0 = Favorable). Error bars indicate 95% confidence interval of the mean. Control indicates model results without injection of tolerance. Other categories correspond to the values of tolerance injected into the 10 most important nodes from 0.6 (TOL60) to 0.0 (TOL0) as ranked by betweenness centrality.

Our results indicate that adjusting the tolerance threshold for the six nodes with the highest betweenness ranking (2.4% of the network) can have a dramatic effect on the ability of an exogenous media campaign to shape the opinions and behaviors in the network. Raising the tolerance value of those six nodes to 0.6 resulted in an increase in the number of smokers from approximately 68% to approximately 90%. The average opinion showed a similar increase from approximately 0.75 to 0.92. Lowering the tolerance threshold for those six nodes strongly mitigated the ability for the media node to influence the network. With no educational or counter-marketing campaigns, decreasing the tolerance of the six highest betweenness nodes reduces the average opinion and smoker fractions toward baseline levels steadily, culminating in the lowest set of values when tolerance is set to 0 which corresponds to no opinion propagation across the six nodes.

6.3 Effects of Addiction

Above, we outlined a mapping between the opinion of an individual and their behavior using step functions at the initiation threshold and at the cessation threshold. The behavioral function takes on the values [0, 1], equivalent to a false/true distinction when asking if the individual engages in the given behavior. We use a value for an initiation threshold, such that an opinion below the threshold value results in no change, while an opinion equal to or above the threshold value results in the individual initiating the behavior.

The initiation threshold may be interpreted as the minimum value an individual's opinion needs to be in order to choose to assume the costs involved in the behavior. Cost here refers to both the direct economic costs, as well as the cost in time and effort. The initiation threshold might be raised by raising the purchase price on the item, or by making the item harder to acquire or consume. Concerns about health effects and addiction would not affect the initiation threshold in this model, but would rather be seen as acting to lower an individual's opinion about smoking.

The effects of addiction are implemented with the introduction of a cessation threshold. The cessation threshold may be equal to or less than the initiation threshold. If the cessation threshold is less than the initiation threshold, this indicates that the opinion of the individual needs to fall lower than it would otherwise, due to addiction acting as an additional motivating component. Thus, for a strongly addictive product, the cessation threshold could be set at 0.35, versus an initiation threshold of 0.65, incorporating the fact that, once an individual is addicted, their ability to quit is compromised – their opinion of the product might fall well below the initiation threshold, but they will continue its use (Figure 6). Strategies that would make it easier for people to overcome the effects of addiction, such as increasing the availability of nicotine replacement therapy or of smoking cessation counseling, would change (raise) the cessation threshold, making it possible for people to quit smoking more easily (at a higher opinion threshold).

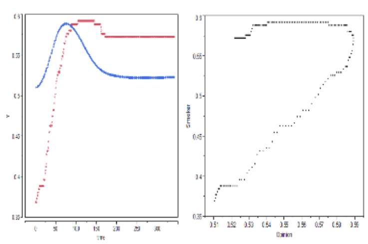

Figure 6: Left diagram shows the effect of advertising and educational campaigns on opinion towards smoking (shown in blue) and smoking behavior (shown in red). Advertising is active through t=75. Education campaign starts at t=75, and induces a decrease in average opinion, but has less of an effect on smoking behavior due to hysteresis. Right diagram illustrates percentage of smokers versus average opinion. The number of smokers increases roughly linearly with rising average opinion, but remains fairly level after the average opinion starts to fall.

7 Conclusions

These analyses demonstrate the value of simple social-network concepts in addressing prevention and treatment of a chronic disease with behavioral components. We have shown a plausible mechanism for pro- or anti-tobacco messages to shift the opinions of a population relative to tobacco and eventually to affect the proportion of individuals who smoke. Using our opinion dynamics model, we have shown the effect of targeted advertising or educational campaigns where message recipients are selected by their network characteristics such as PageRank and betweenness centralities. We have explored two approaches to imposing changes onto a social network, either through the informational content of the message or through enabling better information propagation through the network. The symbiotic effects of a mixed-message advertising campaign was described, showing how a moderate campaign can be applied to increase the effectiveness of a more extreme campaign. Lastly, we developed a straightforward network model of addiction that demonstrates notional match with observed metrics.

In general, the agent-based model is intended to consider the effects of interventions as general characterizations, rather than exactly replicating historical data. By looking at dynamics across over a hundred thousand randomly generated networks with randomly generated initial distributions of opinions and behaviors, we can test interventions for robustness across a wide range of different communities and discover the key components for creating robust interventions.

Bibliography

[1] Albert, R., & Barabási, A. L. (2002). Statistical mechanics of complex networks. *Reviews of modern physics*, *74*(1), 47–97.

[2] Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509.

[3] Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer networks and ISDN systems*, *30*(1-7), 107–117.

[4] Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of modern physics*, *81*(2), 591–646.

[5] Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, *357*(4), 370.

[6] Christakis, N. A., & Fowler, J. H. (2008). The collective dynamics of smoking in a large social network. *New England journal of medicine*, *358*(21), 2249.

[7] Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, *3*(1-4).

[8] Dorwin Cartwright, F. H. (1956). Structural Balance: A Generalization of Heider's Theory. *Psychological Review*, *63*(5).

[9] Farrelly, M. C., Davis, K. C., Haviland, M. L., Messeri, P., & Healton, C. G. (2005). Evidence of a Dose–Response Relationship Between" truth" Antismoking Ads and Youth Smoking Prevalence. *American journal of public health*, *95*(3), 425.

[10] Galea, S., Nandi, A., & Vlahov, D. (2004). The social epidemiology of substance use. *Epidemiologic reviews*, *26*(1), 36.

[11] Glass, L. M., & Glass, R. J. (2008). Social contact networks for the spread of pandemic influenza in children and teenagers. *BMC Public Health*, *8*(1), 61.

[12] Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artifical Societies and Social Simulation (JASSS) vol*, *5*(3).

[13] Hjermann, I., Holme, I., Byre, K. V., & Leren, P. (1981). Effect of Diet and Smoking Intervention on the Incidence of Coronary Heart Disease:: Report from the

Oslo Study Group of a Randomised Trial in Healthy Men. *The Lancet*, *318*(8259), 1303–1310.

[14] Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.

[15] Jager, W., & Amblard, F. (2005). Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change. *Computational & Mathematical Organization Theory*, *10*(4), 295–303.

[16] Jarvis, M. J. (2004). Why people smoke. *British Medical Journal*, *328*(7434), 277.

[17] Ling, P. M., & Glantz, S. A. (2002). Why and how the tobacco industry sells cigarettes to young adults: evidence from industry documents. *American Journal of Public Health*, *92*(6), 908.

[18] Mokdad, A. H., Marks, J. S., MP, M., & Stroup, D. F. (2004). Actual Causes of Death in the United States, 20Ò0. *JAMA*, *291*(10).

[19] Pierce, J. P., Choi, W. S., Gilpin, E. A., Farkas, A. J., & Berry, C. C. (1998). Tobacco industry promotion of cigarettes and adolescent smoking. *JAMA: the journal of the American Medical Association*, *279*(7), 511.

[20] Rigotti, N. A. (2002). Treatment of tobacco use and dependence. *New England Journal of Medicine*, *346*(7), 506–512.

[21] Rosenquist, J. N., Murabito, J., Fowler, J. H., & Christakis, N. A. (2010). The Spread of Alcohol Consumption Behavior in a Large Social Network. *Annals of Internal Medicine*, *152*(7), 426.

[22] Scott, J. (2000). *Social Network Analysis: A Handbook*. London: Sage Publications Ltd.

[23] Single, E., Rehm, J., Robson, L., & Van Truong, M. (2000). The relative risks and etiologic fractions of different causes of death and disease attributable to alcohol, tobacco and illicit drug use in Canada. *Canadian Medical Association Journal*, *162*(12), 1669.

[24] Smith, K. P., & Christakis, N. A. (2008). Social networks and health. *Annual review of sociology*, *34*, 405-429.

[25] Stampfer, M., Hu, F., Manson, J., Rimm, E., & Willett, W. (2000). Primary prevention of coronary heart disease in women through diet and lifestyle. *New England Journal of Medicine*, *343*(1), 16-22.

[26] Valente, T. W. (2003). Social network influences on adolescent substance use: An introduction. *Connections*, *25*(2), 11–16.

[27] Weisbuch, G., Deffuant, G., Amblard, F., & Nadal, J. P. (2002). Meet, discuss, and segregate! *Complexity*, *7*(3), 55–63.