

An opinion-driven behavioral dynamics model for addictive behaviors

Thomas W. Moore^{1,a}, Patrick D. Finley¹, Benjamin J. Apelberg², Bridget K. Ambrose², Nancy S. Brodsky¹, Theresa J. Brown¹, Corinne Husten², and Robert J. Glass¹

¹ Sandia National Laboratories, P.O. 5800 MS1138, Albuquerque, NM 87185-1138, USA

² Center for Tobacco Products, U.S. Food and Drug Administration, 9200 Corporate Blvd. (Room 120M), Rockville, MD 20850, USA

Received 2 September 2014 / Received in final form 15 January 2015

Published online 8 April 2015

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract. We present a model of behavioral dynamics that combines a social network-based opinion dynamics model with behavioral mapping. The behavioral component is discrete and history-dependent to represent situations in which an individual's behavior is initially driven by opinion and later constrained by physiological or psychological conditions that serve to maintain the behavior. Individuals are modeled as nodes in a social network connected by directed edges. Parameter sweeps illustrate model behavior and the effects of individual parameters and parameter interactions on model results. Mapping a continuous opinion variable into a discrete behavioral space induces clustering on directed networks. Clusters provide targets of opportunity for influencing the network state; however, the smaller the network the greater the stochasticity and potential variability in outcomes. This has implications both for behaviors that are influenced by close relationships versus those influenced by societal norms and for the effectiveness of strategies for influencing those behaviors.

1 Introduction

The investigation of health-related behaviors using social network analysis and modeling is a growing field that integrates behavioral, epidemiological, and computational research, and often utilizes models and methods of analysis from statistical physics. Although social network analysis has existed as a research tool for almost a century, a series of recent innovations has advanced this vital and growing area of investigation. Some innovations integrate computational modeling of social processes with an emphasis on explanation or prediction of social phenomena, including the network-scale results of the diffusion of ideas and individual opinion formation. Many topics of interest to social scientists and epidemiologists, however, center on social influences affecting individual risk behaviors.

The domain of sociophysics, the investigation of social phenomena using relatively simple computational models of stylized interpersonal interactions drawn from models developed for statistical physics, has foundations in the extensive work of Serge Galam [1–3]. Galam has led the investigation of the application of these models to a wide range of social and sociopolitical phenomena [4–6]. One broad field of particular interest is the modeling of social influences on individual opinions and behaviors, broadly termed opinion dynamics modeling [1,7,8].

Opinion dynamics models are quantitative computational techniques for modeling the propagation of opinions between pairs of individuals. This class of models derives from the Ising spin models employed in statistical physics, taking from those models the notion that the state of one element in a collection can, through interaction, affect the state of another [1,9]. Although these models are highly simplified implementations of extremely complex psychosocial phenomena, they have generated a great deal of interest in the field of computational social sciences due to their computational and analytical tractability.

Broadly, opinion dynamics models capture the social components of opinion formation as defined by structural balance theory. First proposed in 1953, structural balance theory suggests that individuals connected by positive affect will influence each other in their opinions regarding a third individual or idea [1,10]. This theoretical construct is similar to a model proposed by French [11,12], and to phenomena observed in the famous Asch experiments in social influence in which subjects altered their answers about the perceived lengths of lines in order to conform to group opinions [13,14]. Recognizing that individuals obtain and contextualize information largely from external sources, modeling that is inclusive of these externalities conforms to current thinking in social psychology and provides the flexibility to model diverse phenomena including the effects of different classes of interventions to influence opinions. This approach is an extension of research models that

^a e-mail: tmoore@sandia.gov

use simple or complex propagation of behavioral states, and can suggest additional metrics for characterizing the propagation of behavioral phenomena.

Opinion dynamics models have been applied to several domains, including the spread of environmentally-aware agricultural practices among European agricultural producers and the formation of extremist groups in social networks [1–3]. These models have taken several forms, including: discrete and continuous models of opinion space [4–6]; well-mixed and network-based population interactions [1,7,8]; and various models of opinion propagation including neighbor sampling [1,9], summation [1,10], averaging [11,12], and Bayesian estimation [13].

The influence of social factors on health-related behaviors have generated considerable interest among researchers; social network effects on smoking, obesity, and other health-related behaviors and conditions have been identified across a range of datasets [15–17]. These behaviors and conditions demonstrate assortativity in networks creating clusters of similarly characterized individuals.

In this paper, we describe a behavioral dynamics model that combines opinion dynamics in directed graph social networks, conflicting media messages, opinion-driven behavior and the effects of addiction. Section 2 outlines the conceptual and mathematical components of the model using smoking as the motivating example. Section 3 reports the results of the model including the effects of each input parameter on model behavior. Section 4 concludes with a summary of our findings.

2 Theory and model development

2.1 Opinion

Opinion dynamics models incorporate individual agents to simulate social influences on individual opinion formation. These agents are characterized with a numeric opinion value, which is often implemented using a continuous variable. Opinions are influenced by other individuals in the community, either by random sampling in a well-mixed population or as constrained by a social network topology. In bounded confidence models individuals are also characterized by a tolerance value. This tolerance value contributes a homophily-like characteristic that constrains opinion-changing interactions to individuals exhibiting similarity in opinion, where individual tolerance determines the necessary degree of similarity. For a comprehensive review of different approaches to opinion dynamics modeling, see Castellano et al. [9].

2.1.1 The meaning of opinion

As an illustrative example, we consider opinions about smoking and their influence on smoking behavior. In this model, opinion is represented as an aggregated value summing both utilitarian and affective assessments of smoking. Positive evaluations about smoking tend to increase

opinion about smoking; negative evaluations tend to decrease it. Individuals form opinions on the basis of interactions with their peers and other individuals in their social networks, as well as through the influence of marketing or public health education campaigns.

Utility evaluations are primarily pragmatic: positive utility might be perceived if an individual believes that smoking helps with weight control or stress relief, while negative utility would arise from awareness of the strong causal link between smoking and multiple diseases. Affective evaluations are more image-centered and emotional; the perception of smoking as an assertion of maturity and independence would contribute positively to opinion, while the perception of smoking as associated with undesirable or anti-social behaviors would contribute negatively. There is a large body of empirical evidence identifying the correlation of positive utility and social perceptions with a propensity among youth to smoke, while negative social perceptions and concerns about health consequences are correlated with a decreased likelihood of youth smoking [18].

We initialize initial opinion using a continuous variable on the interval $[0, 1]$. With relatively high mixing of opinions among agents and with initial opinions drawn from a uniform random distribution on the interval $[0, 1]$, the consensus opinion tends to converge to a neutral value, at or near 0.5. In the model, opinions approaching 0.0 are more negative towards smoking; those approaching 1.0 are more positive.

2.1.2 Opinion propagation

As in all opinion dynamics models, the opinions of individuals are influenced by their neighbors. We restrict influence on an individual to neighbors in a social network with directed edges. Directionality in friendship relationships has been indicated as important in the spread of smoking-related behaviors, as well as the spread of other behaviors and opinions [15,16].

The continuous-valued opinion dynamics algorithm proposed by Weisbuch and co-workers uses a model of non-directed interactions between agents in a well-mixed (non-networked) population [19]. Agents x and x' selected at random from the population whose opinions differ by less than some value d , would mutually update their opinions according to the equation:

$$\begin{aligned}x &= x + \mu(x' - x) \\x &= x + \mu(x' - x)\end{aligned}\tag{1}$$

for some value μ , a convergence parameter. Weisbuch also investigated the application of the model to undirected scale-free and lattice network topologies [8], but did not discuss clustering effects in terms of network topology-based clusters.

We adapt this model with the following realization in which both directionality imposed by network topology

and averaging over out-degree neighbors are applied

$$x_i(t+1) = x_i(t) + \frac{1}{|S_i|} \sum_{j \in S_i} \mu_{ij} [x_j(t) - x_i(t)]. \quad (2)$$

In equation (2), the subscript i represents an individual, subscript j is a neighbor to that individual, $x_i(t+1)$ is the opinion value of individual i at the next time step, $x_i(t)$ is i 's opinion at the current time step, S_i is the set of all out-degree neighbors of i and $|S_i|$ is the cardinality of S_i , μ_{ij} is the plasticity value associated with the relationship between individual i and neighbor j , and $x_j(t)$ is the opinion of neighbor j at the current time step.

The averaging function is similar to that proposed by Hegselmann and Krause [11,20] and captures longer-term continuous influences on opinion compared to the randomized discrete interactions utilized in the Weisbuch and co-workers implementation.

The plasticity value functions analogously to an edge weight. Larger plasticity values proportionally increase the influence of the given node. Some empirical research suggests that the quality of a relationship (best friend versus friend, work associate versus spouse) can cause a differential effect in behavior propagation [15,21] thus networks can be constructed using heterogeneous plasticity values to represent different types of relationships.

As in the Weisbuch and co-workers approach, we apply the restriction imposed by tolerance bounds [1]:

$$|x_i(t) - x_j(t)| \leq \varepsilon. \quad (3)$$

In equation (3), ε represents tolerance and indicates the range of opinion to which an individual might be receptive. An individual with higher tolerance will be more open to influence from his or her neighbors, while an individual with lower tolerance will be less open to influence. The number of opinion clusters expected to form in a network is determined by tolerance values [8]. Networks characterized by low tolerance will form many small clusters, while networks with high tolerance will generally come to a single large-cluster consensus value. Network connectivity values also play a role in cluster size distribution. In a well-mixed population with random bi-directional interactions, the number of opinion clusters at steady state will be proportional to $\frac{1}{2\varepsilon}$, with a phase transition from discrete clusters to community-wide consensus at $\varepsilon \approx 0.27$ [6].

2.1.3 Media influences

In this model, opinion-broadcasting media campaigns are represented as nodes with an out-degree of 0; such nodes have the potential to influence individuals but are not themselves subject to influence. The in-degree of a media node is analogous to the reach of the campaign.

In the tobacco example, media sources that have been shown to influence individuals' opinions regarding smoking include cigarette advertising (e.g., displays at points-of-sale), characters smoking in movies, tobacco industry promotions, and direct mail; these sources are

modeled as media nodes with high opinions about smoking [18,22]. Alternatively, public health campaigns and counter-marketing media sources are represented by media nodes with low opinions. A recent study concluded that tobacco control campaigns are more effective in triggering quitting behaviors than advertising campaigns for nicotine replacement therapies (NRTs) [23]. In the context of opinion dynamics, this finding could be interpreted as low opinion about smoking having a greater behavioral effect than a high opinion toward specific NRTs.

Media nodes can be given a lower (or higher) plasticity to reduce (or enhance) their influence relative to that of an individual's friends or others in the network. Media nodes can also be assigned different patterns of connectivity to enable modeling of the effects of different metrics of reach. The influence effects of media nodes are largely a function of message strength, of reach, and of the network centrality of the individuals being influenced [24].

This approach to adding media nodes to opinion dynamics models is similar in intent to Carletti et al. [25], but uses a social network structure rather than a well-mixed population. Media campaigns in this model integrate with the social network topology through connection to a subset of nodes rather than by being applied to all nodes simultaneously. The campaigns are also constantly active, rather than being activated at set time intervals.

2.2 Behavior

To model social influences on the initiation and cessation of individual behaviors, we add a behavioral component to the individual agents in the model. Although the diffusion of information within social networks is an important phenomenon, it is the behavior of individuals within the network that is of greatest interest to many domains, including public health. Product advertisements, public health campaigns and other forms of media and social influence often seek to reinforce or to change behaviors by providing information, evoking emotions, or otherwise influencing opinions in conscious and unconscious ways [24]. Models of social influence on individual behaviors, including health-related behaviors, should therefore incorporate social influences on opinions.

Although opinions in this model occur in a continuous space, behaviors are modeled as discrete variables. For purposes of simplicity, we introduce a simple step function representing behavior as a binary-valued state variable. An individual may either be engaging in the behavior ($b = 1$) or not ($b = 0$). Although some other models of smoking initiation break the initiation process into multiple stages, our model's step function correlates with smoking surveys in which a simplified behavioral definition is used to categorize an individual as a smoker (see, for example, Alexander et al. [21], where smoking status is based on whether an individual has smoked in the past 30 days; the same definition is used in the National Youth Tobacco Survey [26]). The behavioral function can take on multiple stages or states, as needed.

2.2.1 Initiation

If an individual's opinion regarding a behavior passes an initiation threshold, that individual will initiate the behavior, changing the state variable from 0 to 1. The initiation threshold can be thought of as the threshold established by the costs associated with the behavior. In this context, cost should be understood as a total cost, rather than strictly a monetary cost. Using the example of smoking, cost includes the price of a pack of cigarettes, as well as the convenience costs associated with obtaining cigarettes and their use. Thus, the initiation threshold would be raised by increasing the price of a pack of cigarettes as well as by restricting where smoking can occur, and in the case of adolescent smokers, increased difficulty in acquiring cigarettes and fear of punishment for smoking.

In the tobacco example presented here, costs should not be interpreted as subsuming health-related costs. The perceived impacts to health (whether personal through the direct health consequences of smoking, or secondary due to the impact of second-hand smoke on others) are part of an individual's opinion about smoking. Likewise, other social costs, such as others' dislike for the smell of smoke or concern over the social perception of smokers, are part of the individual's overall opinion of tobacco.

By way of example, we consider an individual with an opinion of 0.70 to have a modestly positive opinion about smoking. This individual considers smoking to be moderately attractive, possibly by associating it with positive imagery ("coolness"), or because of perceived benefits (e.g., weight control). For example, if the initiation threshold is 0.70 or less, the individual in this example will start smoking. If, however, the initiation threshold is 0.75, the individual will not initiate, because her opinion is not high enough to surpass the costs associated with the behavior. We investigate the results of different initiation thresholds at the network level in Section 3.2.8.

Raising the initiation threshold means that an individual's opinion about smoking needs to rise to a higher value before that individual will initiate the behavior. If cigarettes become more expensive or difficult to obtain (for adolescents), or if smoking becomes more inconvenient, an individual will need a commensurably higher opinion about smoking to initiate the behavior. However, if cigarettes are inexpensive and easy to acquire and use, an individual may start using them more quickly at a significantly lower opinion value. Empirical evidence indicates that price increases have the effect of reducing smoking initiation, as do indoor clean air laws which restrict permissible smoking areas (see, e.g., [18,27]). Community-level interventions to reduce minors' access to tobacco products also can reduce initiation by raising the initiation threshold. In this approach, effective public health media campaigns operate by reducing opinions about smoking rather than by increasing the initiation threshold [18,28].

Modeling behavior as a binary realization of opinion was initially proposed in a model developed by Martins [13]. The Martins model uses an explicit Bayesian function in which an individual infers the opinions of

neighbors based on observed (binary) actions of those neighbors. In contrast, our model uses a binary representation of behavior which is influenced by opinion dynamics.

2.2.2 Cessation and addiction

If the behavior being modeled has a physiologically or psychologically addictive component, the direct influence of opinion on behavior can be modified. Although opinion might contribute to initiation and early usage, opinion and behavior can diverge, allowing behaviors to persist despite a declining opinion.

In the case of smoking, for example, research shows that most current smokers indicate that they would like to quit [29], or state that they would not have become smokers if they could make the choice again [30]. Despite the fact that more than half of current smokers make a quit attempt in a given year, only a small fraction are successful [29]. Nicotine, a constituent of tobacco, is highly addictive and the principal compound responsible for physiological addiction to tobacco [31–33].

Addictive effects can be modeled by introducing hysteresis into the opinion-to-behavior mapping function. An individual's opinion may fall below the initiation threshold after smoking is initiated, but if an individual has become addicted, a lower opinion may not be sufficient to instigate cessation.

We implement this behavioral latching through the use of an addiction factor which induces a cessation threshold that is equal to or lower than the initiation threshold. The addiction factor is the difference between the initiation and cessation thresholds and represents the degree of addiction of the individual. The degree of addiction can vary among individuals and the risk of addiction increases the earlier in life that smoking is initiated [31,34]. However, in the results presented here, the addiction factor is set to be uniform across the networks to keep the model simple and the results more easily interpretable.

Figure 1 illustrates the effects of this hysteresis loop on behavior in a hypothetical social network that includes a media node. In this scenario, the opinion of the attached media node is set to a value above a smoking initiation threshold, resulting in an increase in network average opinion and smoking prevalence (peaking at $t = 400$). The media node's opinion of smoking is then set to a low value (it becomes a public health education node), resulting in declining average network opinion and decreased prevalence ($t = 400$ through $t = 1000$). However, although the broadcast media node opinion is able to decrease network opinion far below its initial value, smoking prevalence remains well above its initial value, reflecting the effects of hysteresis.

These findings indicate the possibility of a public health campaign successfully changing public opinion about smoking without a correspondingly dramatic impact on behavioral prevalence, pointing to the importance of using metrics measuring changes in both opinion and behavior to judge the effectiveness of campaigns.

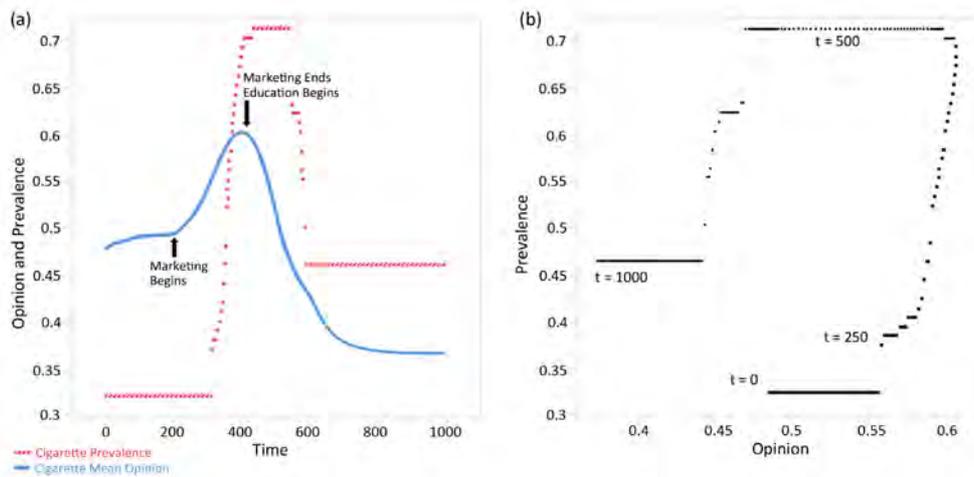


Fig. 1. Plots illustrating the effects of hysteresis. Plot (a) shows a time series of average opinion and prevalence. Plot (b) shows prevalence as a function of average opinion. Plot (b) should be read in a counter-clockwise progression, showing a sharp rise in prevalence as average opinion passes 0.55, and no decrease in prevalence until average opinion falls below 0.45. Prevalence does not decrease to the initial value of approximately 0.325 despite the fact that average opinion falls well below its initial value.

In this example, the cessation threshold could be increased by the introduction of interventions such as raising the costs or barriers associated with smoking. It could also be increased by the use of cessation counseling, the availability of quit lines, or the use of nicotine replacement therapies [35,36]. In this model, these factors effectively reduce the effects of addiction and serve to mitigate barriers to cessation by allowing individuals to quit smoking at relatively higher opinion values.

In domains other than tobacco, the opinion-to-behavior mapping hysteresis function can reflect reluctance to change motivated by sunk costs and loss aversion, by ego-investment, or by physical or physiological constraints.

3 Results

Here we illustrate model functionality. Results presented here derive from simulations run to steady state on directed random graphs of 250 nodes with a Poisson degree distribution (Erdős-Rényi model [37], or ER graphs) using a probability of connection of 0.023 (producing a single component network with an expected density of approximately 0.012 and a mean node degree of 5.75). ER graphs are not generally considered representative of real-world social networks [38]. They are selected for this work to allow the illustration of model dynamics in a base case scenario rather than as further constrained by non-Poisson degree distributions and network community structure.

We follow the general approach in the literature of assigning initial opinions to the nodes from a uniform random distribution on the interval $[0, 1]$. Opinions approaching 0.0 are interpreted as being negative toward smoking, while those approaching 1.0 are interpreted as being positive. Opinions at or near 0.5 are considered essentially neutral. These assessments of “positivity” and

“negativity” are loose conceptual categorizations rather than model components. An equivalent approach might use an interval $[-1, +1]$, in which positive and negative values are more clearly defined. The uniformity of the interval used for the initial opinion distribution should not be read as suggesting that opinions about smoking, or indeed most other subjects, are uniformly distributed in populations. Rather, it is the assignment of a metric where some opinions are more positive than others.

Our results demonstrate dynamics in a highly abstracted, idealized mathematical environment. In order to isolate the effects of specific contributors, we simulate conditions in which media campaigns run in a relatively exclusive manner with at most two competing campaigns. In real-world scenarios, the situation is considerably more complex, with multiple messages being broadcast from multiple sources simultaneously.

3.1 Clustering effects in opinion and behavior space

Behavioral risk factors such as smoking have been shown to cluster in populations. Peer and familial influences are strong risk factors for initiation and maintenance of smoking [18]. The prevalence of smoking is more heavily concentrated in particular subpopulations, such as individuals with low educational attainment and income [39]. Using data from the Framingham Heart Study, Christakis and Fowler demonstrated clustering of smoking on the periphery of a social network and the influence of close network ties’ cessation behavior on an individual’s likelihood of quitting smoking [15].

Mapping a continuous opinion variable into a discrete behavioral space induces clustering on directed networks. Figure 2 illustrates the differences in clustering characteristics in opinion and behavioral spaces as a function of the tolerance parameter.

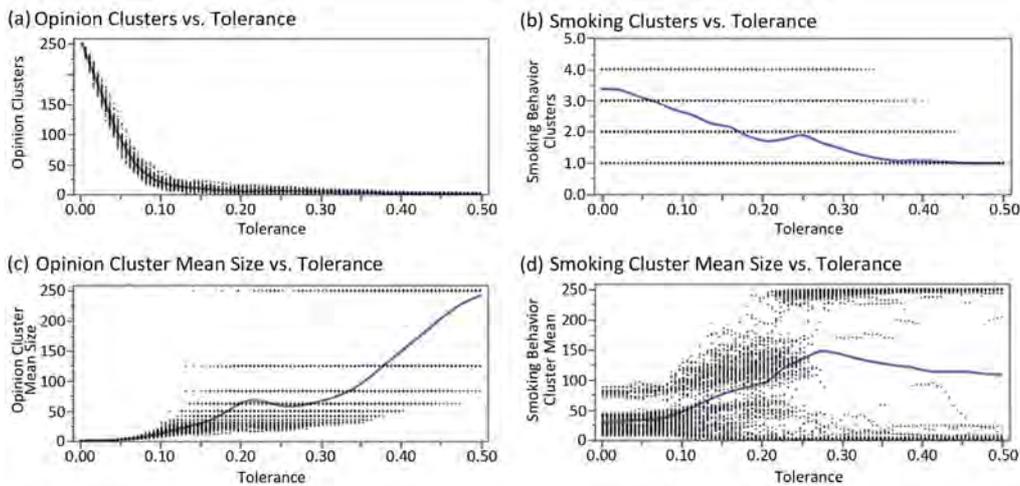


Fig. 2. Clustering in opinion and behavior spaces: points indicate equilibrium values for each particular run, lines show mean values. Opinion clusters represent connected nodes of similar opinion. Two connected nodes are considered to be part of the same opinion cluster if their opinions at steady state differ by less than their tolerance values. Behavior clusters represent clusters of smokers. Results are based on 100 individual runs at each of 100 data points between tolerance values of 0.0 and 0.5, for a total of 10 000 runs.

In this example, two connected nodes are considered to be part of the same opinion cluster if their opinions at steady state differ by less than the tolerance value. Figure 2a illustrates that as tolerance (or the potential to be influenced) increases, the number of opinion clusters decreases dramatically. This reflects the process by which the network achieves a consensus opinion. Behavioral clusters show a similar behavior, but with greater variability and far smaller numbers, as shown in Figure 2b. As the number of opinion clusters decreases, the average size of clusters increases, as shown in Figure 2c. In behavioral space, however, we observe a bifurcation in cluster size. At high tolerances, the network is dominated by either large or small groups of smokers, as shown in Figure 2d.

3.2 Parameter investigation

3.2.1 Model parameter effects

All model parameters are provided in Table 1. By extending the model to incorporate opinion-driven behaviors and addiction, we have increased the number of parameters relative to the cited opinion dynamics models. In this section, we examine model responses for variations across these parameters. In addition to the parameters listed here and investigated below, edge permeability is also considered a parameter. However, edge permeability, if uniform throughout a network, has been shown to control the rate of convergence rather than inform the final state of the model. Because we hold permeability uniform for the scenarios investigated here, we do not include a separate investigation of network edge permeability. The same does not hold true for media node edge permeability, which may be higher or lower than network edge permeability.

To assess the impact of parameters we employ randomly generated networks. Parameter values are held to

the default values given in Table 1 with the exception of the parameter under investigation which is swept through the possible range of values. Each presented scenario used 100 data points sampled uniformly from the distribution of the parameter being investigated. One hundred random networks were used at each parameter value to determine node opinion and network prevalence. The model simulations were considered complete when the opinions and behaviors of the nodes reached equilibrium. This process resulted in 10 000 model runs (100 randomly generated networks and populations \times 100 parameter values) for each parameter investigated. We present results using contour plots of opinion and behavior space. For behavior, the X axis represents the value of the parameter being swept, the Y axis represents the prevalence of smoking at that point, and the color represents the number of networks with that final prevalence value. For opinion space, contour plots of individual nodes are similarly generated, with the X axis representing the parameter under investigation and the Y axis representing opinion where color indicates the number of individual agents sharing that opinion value. For prevalence, each point on the X axis can be thought of as a column, with the sum of the values represented by the colors of the contour map totaling 100 (the number of stochastic runs taken at that X value). For opinion, each column sums to $100N$, where N is the number of nodes used in that scenario. Finally, we use line graphs or other plots to show the mean values of the output variables to assist with interpretation.

3.2.2 Number of nodes

In this section we demonstrate the effects of network size on the behavior of the model. In the present work we are primarily interested in the dynamics of smaller networks,

Table 1. Model parameters.

Parameter name	Possible range of values	Default used in simulations	Discussion of choice for value
Number of nodes (N)	0–2500 Note that this is not a limitation imposed by model architecture; rather, the number of nodes is limited only by the practicality of memory and simulation run times.	250	As shown in Section 3.2.2, average results are relatively insensitive to number of nodes, although differences exist when model response surfaces are viewed in detail. We select a default value of 250 to analyze the dynamics of small networks representative of high school grades, and to avoid under-representing stochastic effects via the reduced variability exhibited by larger ER networks.
Probability of connection (network density)	Depends on number of nodes; ranges from low enough to achieve a single component graph of low density to high enough to produce a fully connected graph of density 1.	0.023	Section 3.2.3 shows that as value increases from 0 to 0.05 the network moves towards a single component. Above 0.05, network behavior is relatively insensitive to this parameter. While the average number of friendship nominations in the National Longitudinal Study of Adolescent Health (Add Health ¹) dataset varies by age and sex, the average number of nominations is 5.9. Although we intentionally are not replicating the community structure of real-world social networks, the expected number of edges in our graph ($0.023 \times 250 = 5.75$) is sufficient to give a connected graph and is in line with the Add Health data.
Media node opinion	0–1	0.375 for Public Health Campaigns, 0.725 for smoking promotions	Section 3.2.4 shows that greatest efficacy in affecting behavior is gained if value is set to just above initiation threshold or just below cessation threshold. For purposes of illustration, we hold Media Node Broadcast Values to within this high efficacy range.
Media node edge count	0– N	25	As shown in Section 3.2.5, the ability of Media Node Edge Count to influence opinion and prevalence saturates by $2N$ for networks with 250 nodes. For purposes of general illustration, we hold Media Node Edge Count to 25 for networks of 250 nodes to give media influence a strong but suboptimal ability to influence the network.
Media node edge permeability	0–1	0.05	As shown in Section 3.2.6, the ability of Media Node Edge Permeability to influence opinions and behaviors saturates at values above twice the default network edge permeability. For purposes of general illustration of media effects, we select 0.05 as a default value to give media influence a strong but suboptimal ability to influence the network.
Tolerance	0–1	0.21	As shown in Section 3.2.7, tolerance values between 0.15 and 0.25 place the network into an influenceable but non-consensus region. We selected a value within this region with a minor bias toward consensus based on an analysis of Figure 3a, which shows 0.21 to be the threshold of maximum influence of the media node.

¹ Add Health is a longitudinal dataset comprising friendships and behaviors for a nationally representative sample. Subjects were in 7th–12th grades during the 1994–1995 school year.

Table 1. Continued.

Parameter name	Possible range of values	Default used in simulations	Discussion of choice for value
Initiation threshold	0–1	0.70	As shown in Section 3.2.8, prevalence is strongly influenced by initiation threshold. For purposes of illustration, we select a threshold that yields prevalences between 30 and 50 percent. Actual smoking prevalences vary among countries and demographic groups, ranging from over 70% of the adult male population in Russia, to less than 10% in Ethiopia [34].
Addiction factor	0–Initiation threshold	0.30	As shown in Section 3.2.9, addiction maintains behavior of smokers, even though their current opinion of smoking may be less than it was when they initiated. We select a baseline value of 0.3 for illustration purposes because it results in a prevalence distribution matching that described in the Initiation Threshold section.

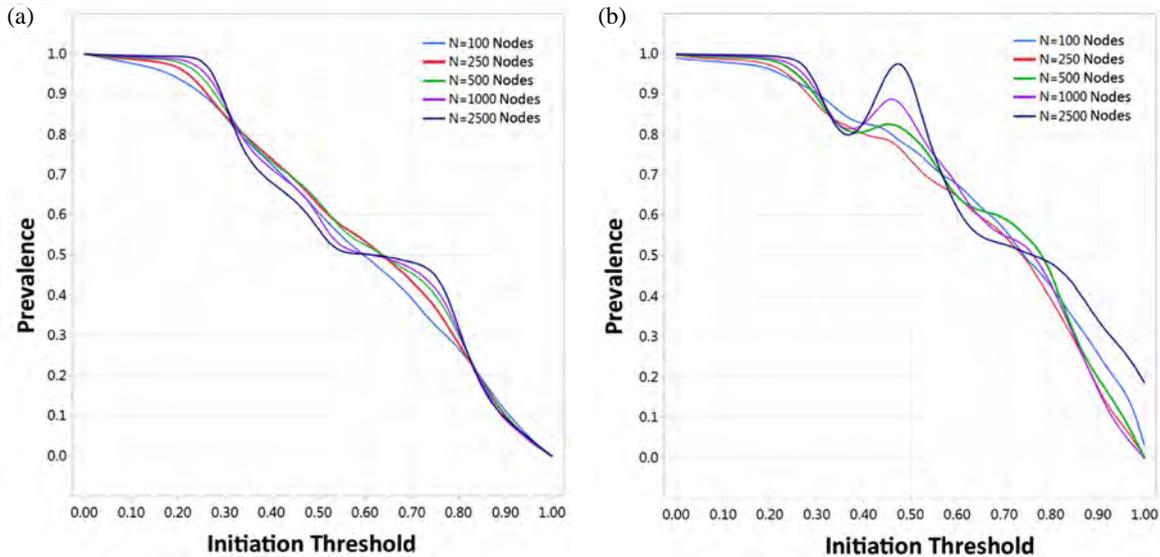


Fig. 3. Comparison of initiation threshold sweep for different network sizes. (a) Average prevalences for networks of 100, 250, 500, 1000, and 2500 nodes. (b) Prevalence in the same networks under the influence of advertising.

and have selected 250 nodes as representative of a network size that might represent the students of a particular grade level in a US high school, or a small network in a larger community consisting of friends and friends of friends. Although the ER topology is non-representative of real-world social networks, smaller network dynamics are relevant and desired. However, in this section we relax the small network constraint and examine the behavior of the model on ER graphs with sizes ranging from 100 to 2500 nodes.

In this section, we compare system responses at 250, 500, 1000, and 2500 nodes, holding network density constant just above the connection threshold of $\ln(N)/N$, for networks with N nodes [35] (probability of connection $p = \ln(N)/(N \times 1.05)$). For purposes of comparison, we look at a scenario in which we sweep the initiation thresh-

old over the interval $[0, 1]$ with both the presence and absence of an advertising node with a broadcast opinion of 0.025 above the initiation threshold and with the number of edges held at 10% of network size. Full analysis of this and contrasting scenarios of networks with the default size of 250 are presented in Section 3.2.8.

In their mean behaviors, networks without media influence behave very similarly independent of size, as shown in Figure 3a. In Figure 3b, we observe that the networks again behave similarly, with the exception of a region between initiation values of 0.4 and 0.55. The reason for this effect can be seen in Figures 4 and 5. Increasing the number of nodes from 100 to 2500 increases consolidation in opinion space from a relatively undifferentiated field at 100 nodes through the formation of upper, lower, and central attractors. The different network sizes shown here

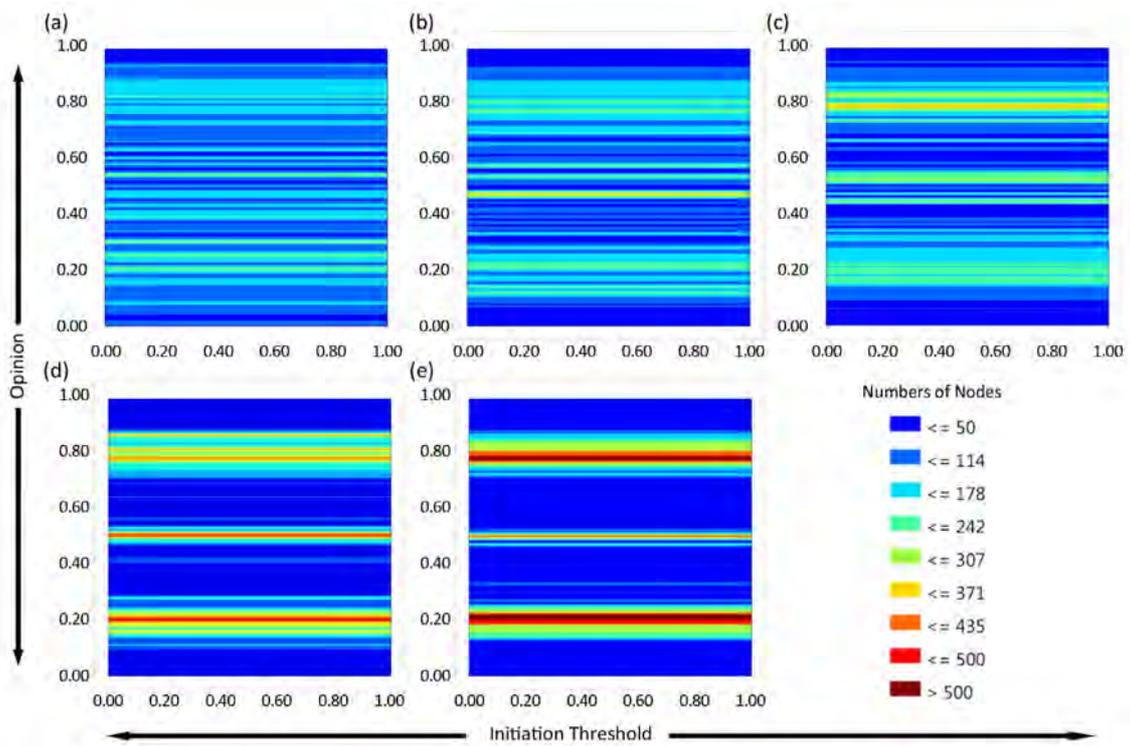


Fig. 4. Opinion space contour maps for different network sizes. (a) 100 nodes; (b) 250 nodes; (c) 500 nodes; (d) 1000 nodes; (e) 2500 nodes. Note that the values for the contour key are for the 100 node contour map. Contour values should be multiplied by the factor increase in N to get the contour values for the other maps.

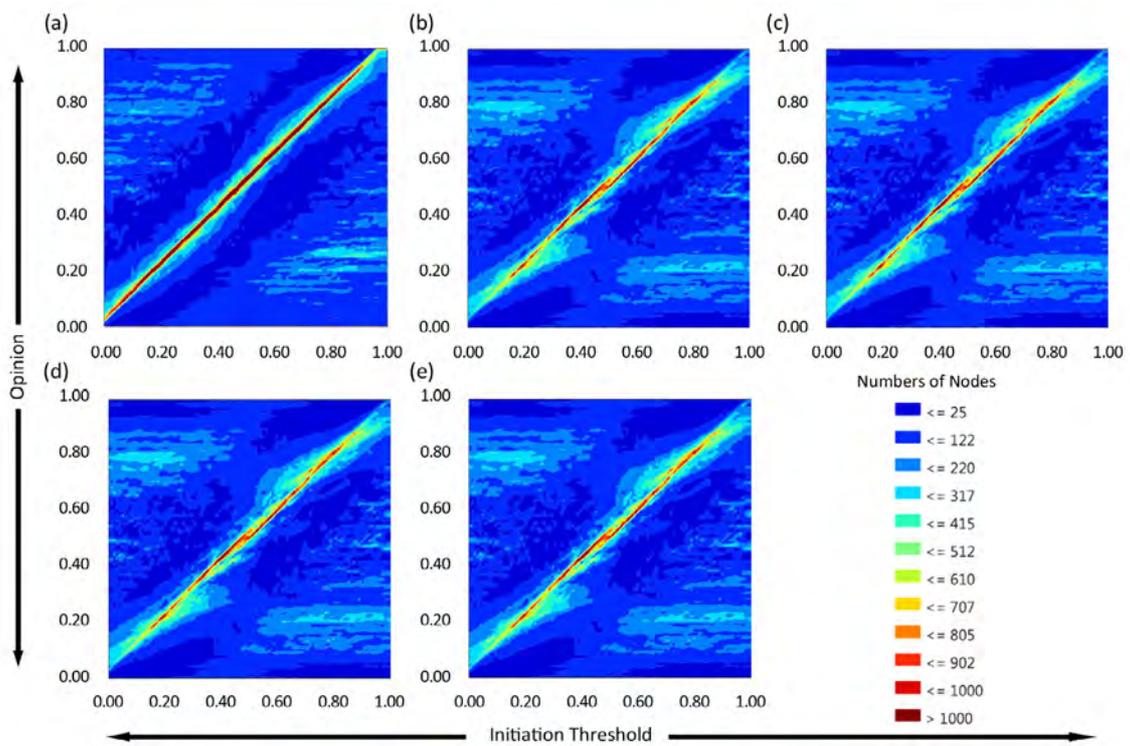


Fig. 5. Opinion space contour maps with media node broadcasting opinion 0.025 above initiation threshold. (a) 100 nodes; (b) 250 nodes; (c) 500 nodes; (d) 1000 nodes; (e) 2500 nodes. Note that the values for the contour key are for the 100 node contour map. Contour values should be multiplied by the factor increase in N to get the contour values for the other maps.

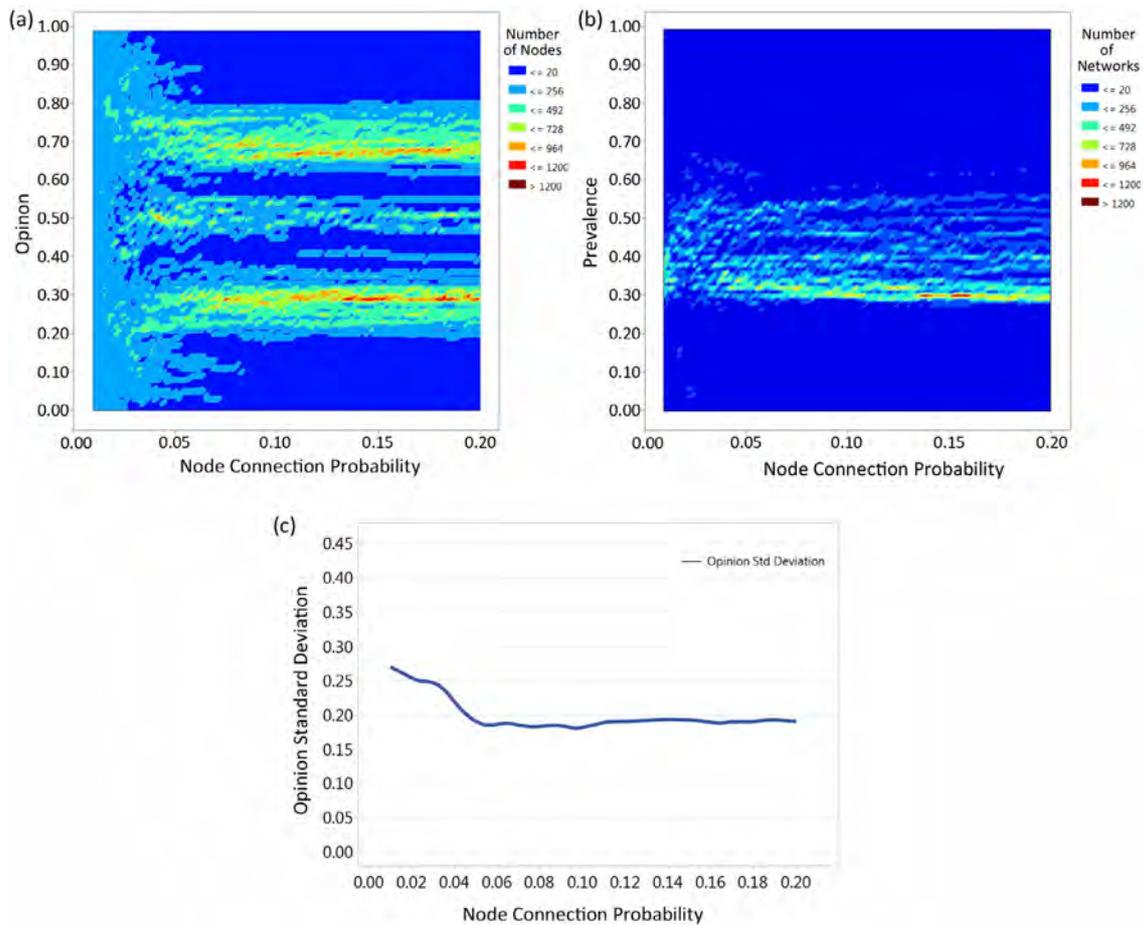


Fig. 6. Opinion (a), prevalence (b) and standard deviation of opinion (c) as a function of node-connection probability. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

result in different banding effects, not necessarily more representative realistic networks. Larger networks could alternatively be achieved by joining clusters together.

These aggregations, or nodes, produce opportunities for influence by the media node. For the high node-count runs, the effect on prevalence is evident when the advertising node moves through these strong bands. The strongest effect is exhibited in the central region, where the initiation threshold is 0.475 and the media node opinion is 0.5. With the ability to draw from increasingly consolidated upper and, more importantly, lower modes, the media node is able to pull a large number of individuals above the initiation threshold values, however, the media node is unable to influence the lower mode, and prevalence levels remain closer to baseline values.

3.2.3 Edge density

Edge density on ER-type graphs is controlled by the probability P that a given node A will be connected to another node B . Holding the number of nodes constant at 250, we

examine how the edge density changes the response of the system.

ER-type graphs approach a single component as P approaches $\ln(N)/N$, with N being the number of nodes in the graph [37]. Below this value, 0.022 in our 250-node example, the network exists with multiple components. Above this value, the network is a single component with network edge density increasing with P .

Without external stimuli such as a media node, edge density has little effect on behavioral prevalence. As shown in Figure 6a, at very low node-connection probabilities (<0.01) the opinion distribution remains uniform, in line with the initial opinion distribution.

Increasing edge density on an ER-type graph induces a three mode distribution with most of the population divided between the upper and lower modes and a lesser number in the middle position. Changing edge density has a negligible effect on prevalence (Fig. 6b) because opinions are consolidated rather than shifted as evidenced by the drop in the standard deviation of opinion with increasing edge density (Fig. 6c).

The addition of an advertising node with 25 edges (10% of the 250-node graph) slightly alters the effects of

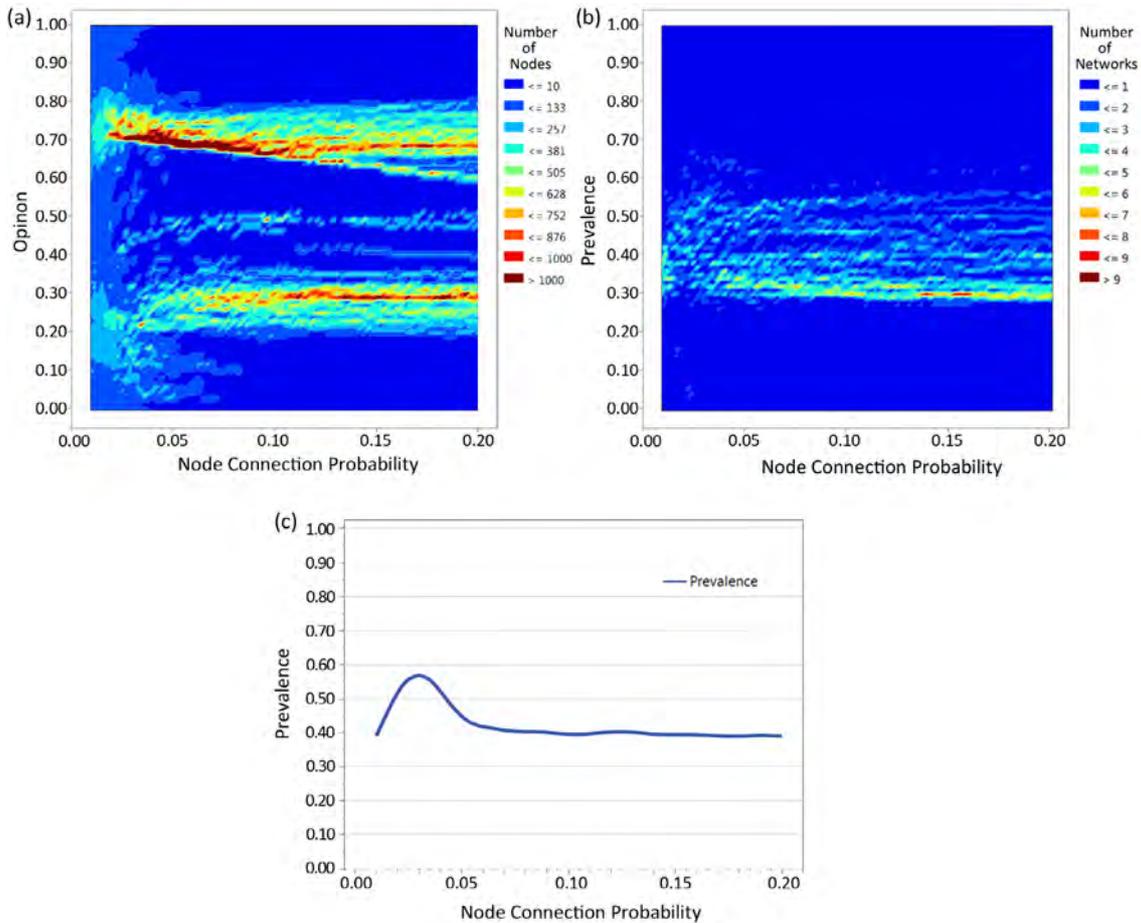


Fig. 7. Opinion (a), prevalence (b) and standard deviation of prevalence (c) as a function of node-connection probability in the presence of a media node broadcasting an opinion of 0.725. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

increasing edge density. Figure 7 shows that the consolidation effect in the higher opinion mode is much stronger, reflecting the influence of the advertising node broadcasting with an opinion value of 0.725. This is also reflected in a corresponding increase in prevalence and the standard deviation for prevalence at node connection probabilities around 0.3. Further increases in the number of edges between network nodes dilute the effects of the advertisement, which drives the lower bound of the upper opinion mode below the initiation threshold of 0.70. A lower initiation threshold would have been able to extend the effective prevalence gain.

Increasing the number of media edges from 25 to 100 to compensate for the dilution effect strongly increases the opinion consolidation effect, as shown in Figure 8a. It also extends the ability of the advertising node to influence behaviors in the network by significantly constraining the formation of a middle mode of opinion. This increase allows the media node to maintain the increase in prevalence over a wider range of edge densities, as shown in Figures 8b and 8c. This indicates that the effectiveness of a media campaign impacting a given number of peo-

ple is proportional to the average number of connections between people connected in the social network.

3.2.4 Media node opinion

The opinion value broadcast by a media node indicates the value to which it may move individuals in the network. The ability of a media campaign node to modify individuals' opinion is limited by individual tolerance bounds. A media node broadcasting an opinion can induce a bimodal distribution in network opinions, with some nodes moving closer to the broadcast opinion while other nodes cluster together indicating the broadcast opinion is outside their tolerance window. Increasing the number or strength of edges serves to consolidate the affected mode rather than to influence individuals to cross from the other mode, as was shown in Figure 8a.

Figures 9a and 9b show the effect of the media node on network opinion and prevalence, respectively, as the media node broadcast opinion is swept from 0 to 1. We observe that, at any given broadcast opinion value, the media node

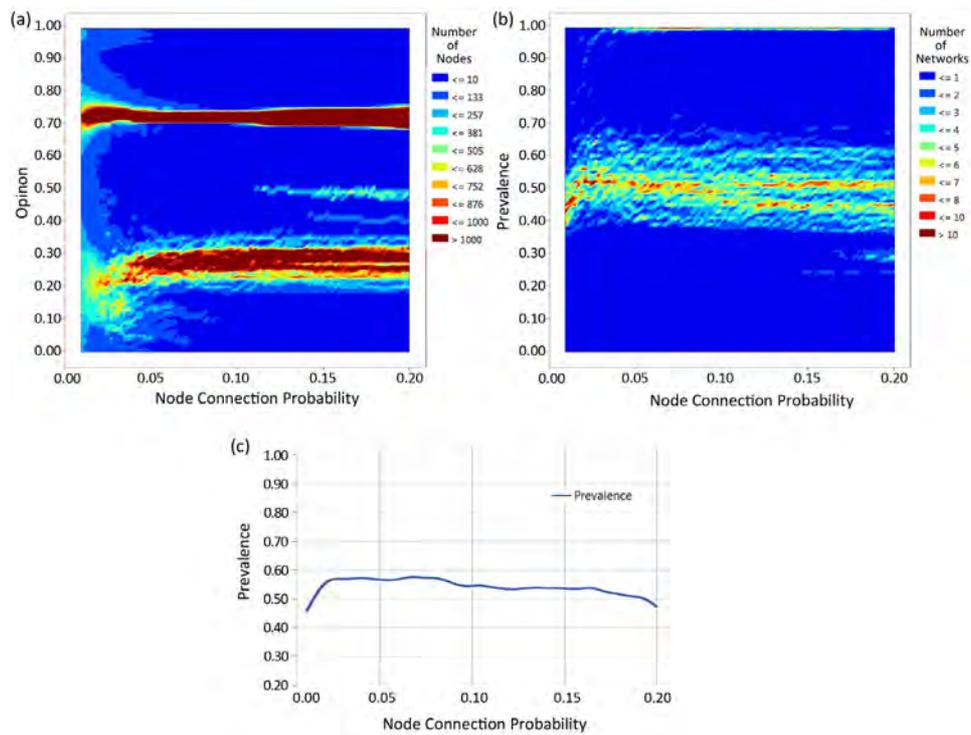


Fig. 8. Opinion (a), prevalence (b) and standard deviation of prevalence (c) as a function of node-connection probability in the presence of a media node broadcasting an opinion of 0.725 with an increased number of media node edges. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

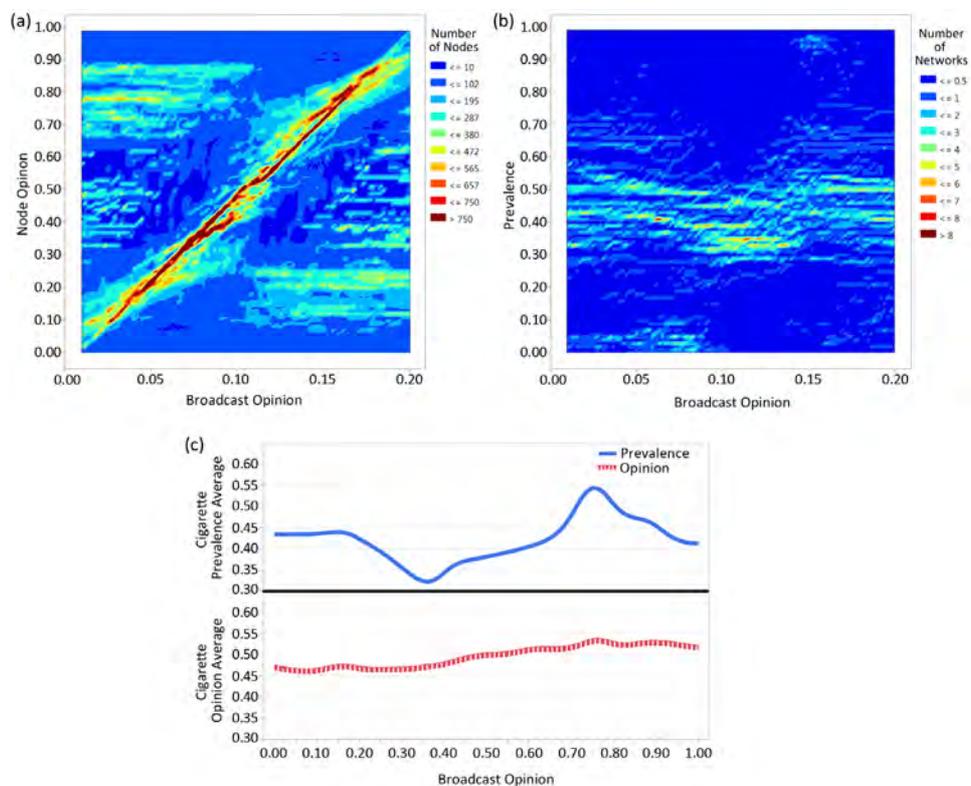


Fig. 9. Opinion (a), prevalence (b) and opinion and prevalence averages (c) as a function of media node broadcast opinion. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

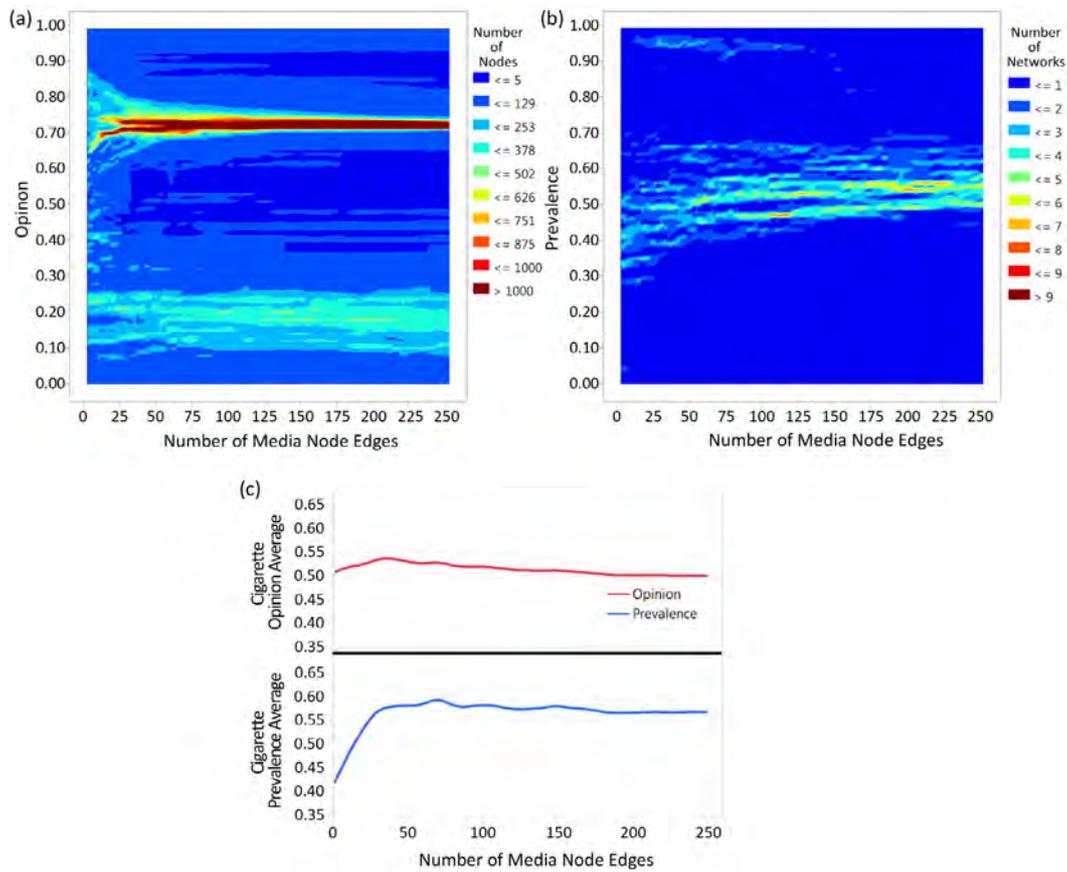


Fig. 10. Opinion (a), prevalence (b) and opinion and prevalence averages (c) as a function of the number of edges for a media node broadcasting an opinion of 0.725. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

is capable of influencing significant numbers of individuals, as shown in the individual opinion contour plot (Fig. 9a). The effect of the broadcast media node on average network opinion (Fig. 9c) is more modest, due to the large number of nodes unaffected by the broadcast opinion due to tolerance constraint. An initiation threshold of 0.7 and an addiction factor of 0.3 (resulting in a cessation threshold of 0.4) are used for calculation of prevalence (Fig. 9b). In driving behavior, we can see in the opinion average plot (Fig. 9c) that the broadcast opinion is most effective at decreasing network opinion when it is slightly below the cessation threshold of 0.4 and most effective at increasing average opinion when it is slightly above the initiation threshold of 0.7, giving the average behavior response a characteristic saw-tooth pattern.

3.2.5 Media edges

Just as the edges between nodes in the model represent the influence of a relationship rather than a single interaction, edges from media nodes represent the aggregated influence of media campaigns rather than the viewing of a single advertisement. The effects of a media campaign

are partially dependent on the number of people reached by the campaign. However, media campaigns are fundamentally constrained by their broadcast opinions and the opinions and tolerances of the nodes in the network.

As we see in Figure 10a, a media node broadcasting an opinion of 0.725 can induce a bifurcation in individual opinions. As the number of media node edges increases to above approximately 10, the opinion distribution becomes bimodal. Further increasing the number of media node edges results in consolidation, or narrowing, of the upper band (Fig. 10a). Increasing the number of edges for this media node up to 30–40 results in an initial marginal increase in average opinion (Fig. 10c) and a stronger increase in prevalence (Figs. 10b and 10c). The increase in prevalence derives from the effect of the media campaign on individuals already near the initiation threshold, as indicated by the contour map (Fig. 10a).

After approximately 30–40 edges have been added, average prevalence stabilizes as shown in Figure 10c. The individuals capable of being influenced by the media campaign, or by their friends who have been influenced, are constrained by tolerance bounds. Additional edges have little additional effect on average prevalence. It can be observed, however, that the consolidating effect in opinion

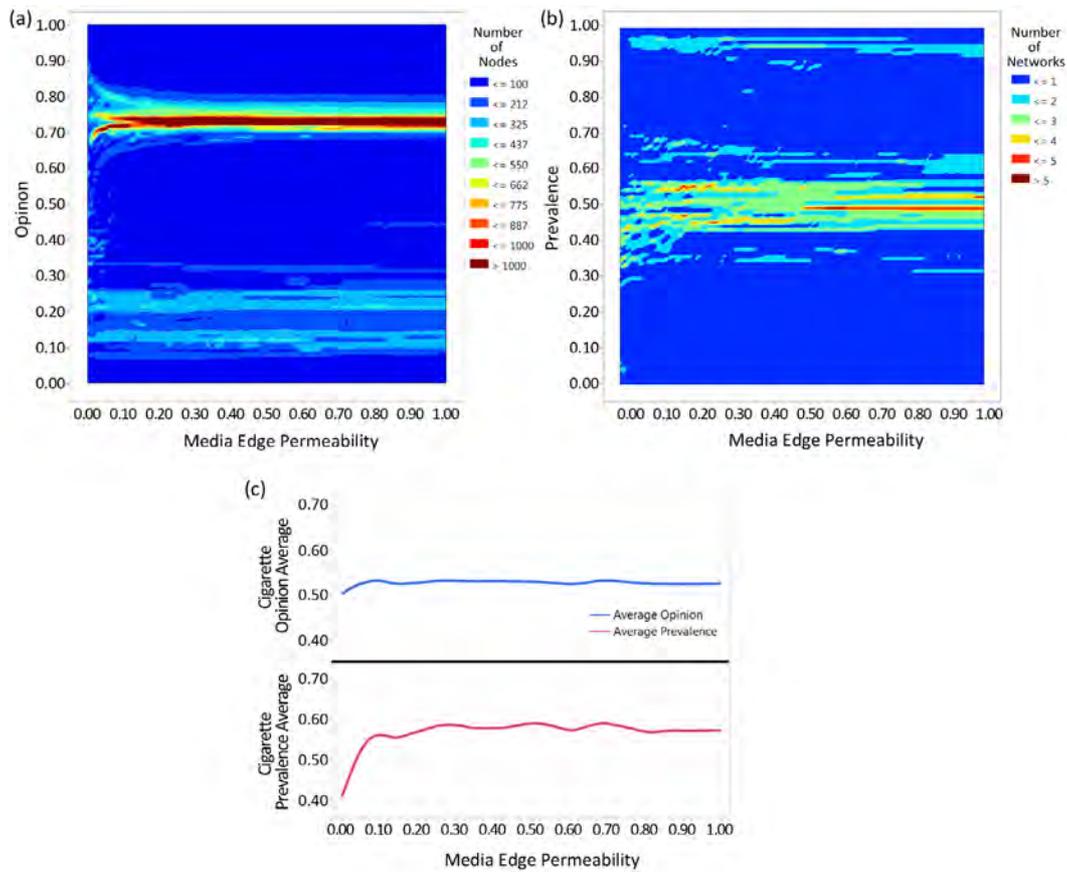


Fig. 11. Opinion (a), prevalence (b) and opinion and prevalence averages (c) as a function of the edge permeability for a media node broadcasting an opinion of 0.725. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

space does decrease outlying results in prevalence as evidenced by the disappearance of the uppermost band of outliers in the prevalence contour plot (Fig. 10b) and the minor decrease in average opinion (Fig. 10c).

3.2.6 Media edge permeability

The permeability, or edge weight, of media nodes can be thought of as the salience of the media campaign to the receiver, assuming that the broadcast opinion is within the receiver's tolerance bounds. In the weighted-average function used to determine an individual's next opinion value, edge permeability contributes a relative weighting, making one individual's opinion more (or less) important than another's. With a baseline permeability of 0.05 throughout the network, a media node with a permeability of 0.025 would contribute exactly half as much to an individual's next opinion value as would a friend whose relationship had the baseline permeability.

The number of media node edges and the media edge permeability are similar in that they affect the ability of the media node to connect into and influence the network, and both are ultimately constrained by the toler-

ance bounds in the targeted network. As shown in Figure 11, the presence of a media node causes a bifurcation in opinion (Fig. 11a), and in prevalence (Fig. 11b). Increasing permeability increases average opinion modestly (Fig. 11c), but has a stronger effect on average prevalence (Fig. 11c). However, after an increase to 0.1 (twice that of the baseline), the effect on prevalence and opinion levels off. As demonstrated by the node opinion contour map (Fig. 11a), this is due to the increase in permeability affecting only those nodes with tolerance bounds that include the media node's broadcast opinion. This consolidates the opinions of individuals within the higher mode band, but does not increase the ability of the media node to recruit individuals from the lower mode band.

3.2.7 Tolerance

Tolerance restricts the range of opinions able to influence a given node. Neighbors with an opinion outside the node's tolerance bounds are ignored. From a network perspective, it is as if the edge connecting them does not exist. This reduces effective network density, but in a dynamic fashion. It also contributes an aspect of dynamic homophily (like assorting with like) to the network.

We observe that, in networks constructed with a node connection probability of 0.023 (yielding a single component sparsely connected with an expected edge density of approximately 0.012), opinions move toward consensus values of 0.5, as expected when drawing initial opinions from a uniform distribution on the interval $[0, 1]$. We note that this convergence occurs at a higher tolerance value (0.5 rather than 0.27) than would be expected in the well-mixed model of Weisbuch et al. [1]. We also note that there remains significant diversity in results throughout the tolerance region $[0.5, 1]$. In a well-mixed population with mutual (non-directed) interactions, the number of clusters formed is approximately the integer part of $1/2d$, where d is the uniform tolerance value [1]. We do observe the formation of some consensus opinion values near 0.2 and 0.8 when tolerance nears 0.2, and the formation of a stronger consensus region when tolerance nears 0.25.

At low values of tolerance (<0.15) opinions remain relatively unchanged from initial conditions (Fig. 12a) but as tolerance increases, opinions are moved toward consensus and most individual opinions converge to 0.5 as tolerance approaches 0.5. Even at a tolerance of 1, which removes the tolerance constraint entirely, heterogeneity in individual opinions can be maintained at steady state due to network connection, directionality and averaging. By way of example, a node A connected to a node B with opinion 0 and a node C at opinion 1 will maintain a steady state opinion value of 0.5, assuming B and C do not reciprocate A 's connections.

The average network opinion (Fig. 12c) is close to 0.5 at low (<0.15) values of tolerance because node opinions were uniformly distributed over the interval $[0, 1]$ and change very little. As tolerance increases, nodes begin to interact, and the range of possible average network opinions increases (Fig. 12c). Correspondingly, networks can stabilize at a wide range of prevalence values (Fig. 12b). As tolerance further increases, consensus forms near 0.5 for most of the networks (Fig. 12c), leaving some outliers due to network topology.

The shift to a wide range of possible prevalence values (Fig. 12b, at tolerance ~ 0.15) precedes the formation of a three mode regime in prevalence in which different stochastic runs resulted in low (0), moderate, (0.3, ± 0.05), or high (1.0) prevalence values. The presence of the significant mode centered on 0.3 is expected due to the initiation threshold of 0.7 and hysteresis value (addiction factor) of 0.3. At the beginning of a scenario run, initial node opinion is drawn from a uniform distribution on $[0, 1]$, resulting in an expected value of 30% prevalence. When, at high tolerance values, nodes approach consensus values near 0.5, individuals who began smoking will generally not achieve an opinion low enough to permit cessation (in this case, an opinion of 0.4 or less). With hysteresis removed, as in Figure 12d, the 0.3-centered mode is absorbed into the lower mode at 0.

The effects of the heterogeneity observed in the opinion space contour map (Fig. 12a) can be seen in plots of the cluster mean size and cluster count (Fig. 12e). We define as a cluster, nodes which are connected and whose

opinions differ by less than 0.02. Even at a tolerance of 1, heterogeneity can be maintained due to network topology and the number of clusters can range up to 10 or higher.

Network density strongly influences the effects of tolerance. By increasing the probability of connection to 0.2, from 0.022, expected network density increases to 0.1. For this higher density network, nodes interact more at lower values of tolerance due to the increase in the number of neighbors and thus the number of expected possible interactions. The formation of a three mode distribution is apparent when tolerance reaches 0.175 (Fig. 13a), with a strong consensus mode forming by 0.25 and bimodality disappearing by 0.3. In a well-mixed population in which interactions are non-directional, the number of modes varies as $1/2d$ where d is the uniform tolerance value in the population [1], predicting a bimodal distribution from 0.175 to 0.25. With directed edges even on a dense ER-type graph, we instead see a three mode distribution through this range, with attractors near 0.3, 0.5, and 0.7. Convergence to opinion consensus at approximately 0.5 does occur at tolerance 0.275, the same as would be expected in a well-mixed network. Prevalence (Fig. 13b) shows a corresponding decrease in variability, with the expected spike in the three mode region induced by the opinion attractor near the initiation threshold of 0.7, and a consolidation above tolerance = 0.3 corresponding to the loss of multimodality.

Tolerance effects under media influence

Tolerance has a strong effect on the ability of a media node to influence the network. We attach a media node with a broadcast opinion of 0.725, using the default parameter values for network generation. Figure 14a shows that tolerances below 0.1 leave the media node unable to effect a change in average network opinion. We observe that weak bimodality in opinion forms around a tolerance of 0.15. A phase transition occurs near tolerance 0.2, induced by the approximate offset between the average network opinion of 0.5 and the broadcast opinion of 0.725, leading to a consolidation around the broadcast opinion for tolerance values above 0.25. As expected, we see a corresponding increase in prevalence (Fig. 14b), with the mode-induced spike at 0.2 followed by a two mode distribution split between full prevalence (1.0) and the hysteresis-induced mode centered around 0.3. The mode at 0 prevalence, as seen in the baseline scenario, has disappeared entirely. As in the previous scenario, however, removal of hysteresis (Fig. 14c) allows the mode clustered at 0.3 to be consolidated at a prevalence of 0.

These results indicate that tolerances in the range of approximately 0.15–0.25 place the system within a critical region of non-consensus but influenceable opinions.

3.2.8 Initiation threshold interventions

Interventions that raise the costs associated with starting to smoke can be modeled as a rise in the initiation

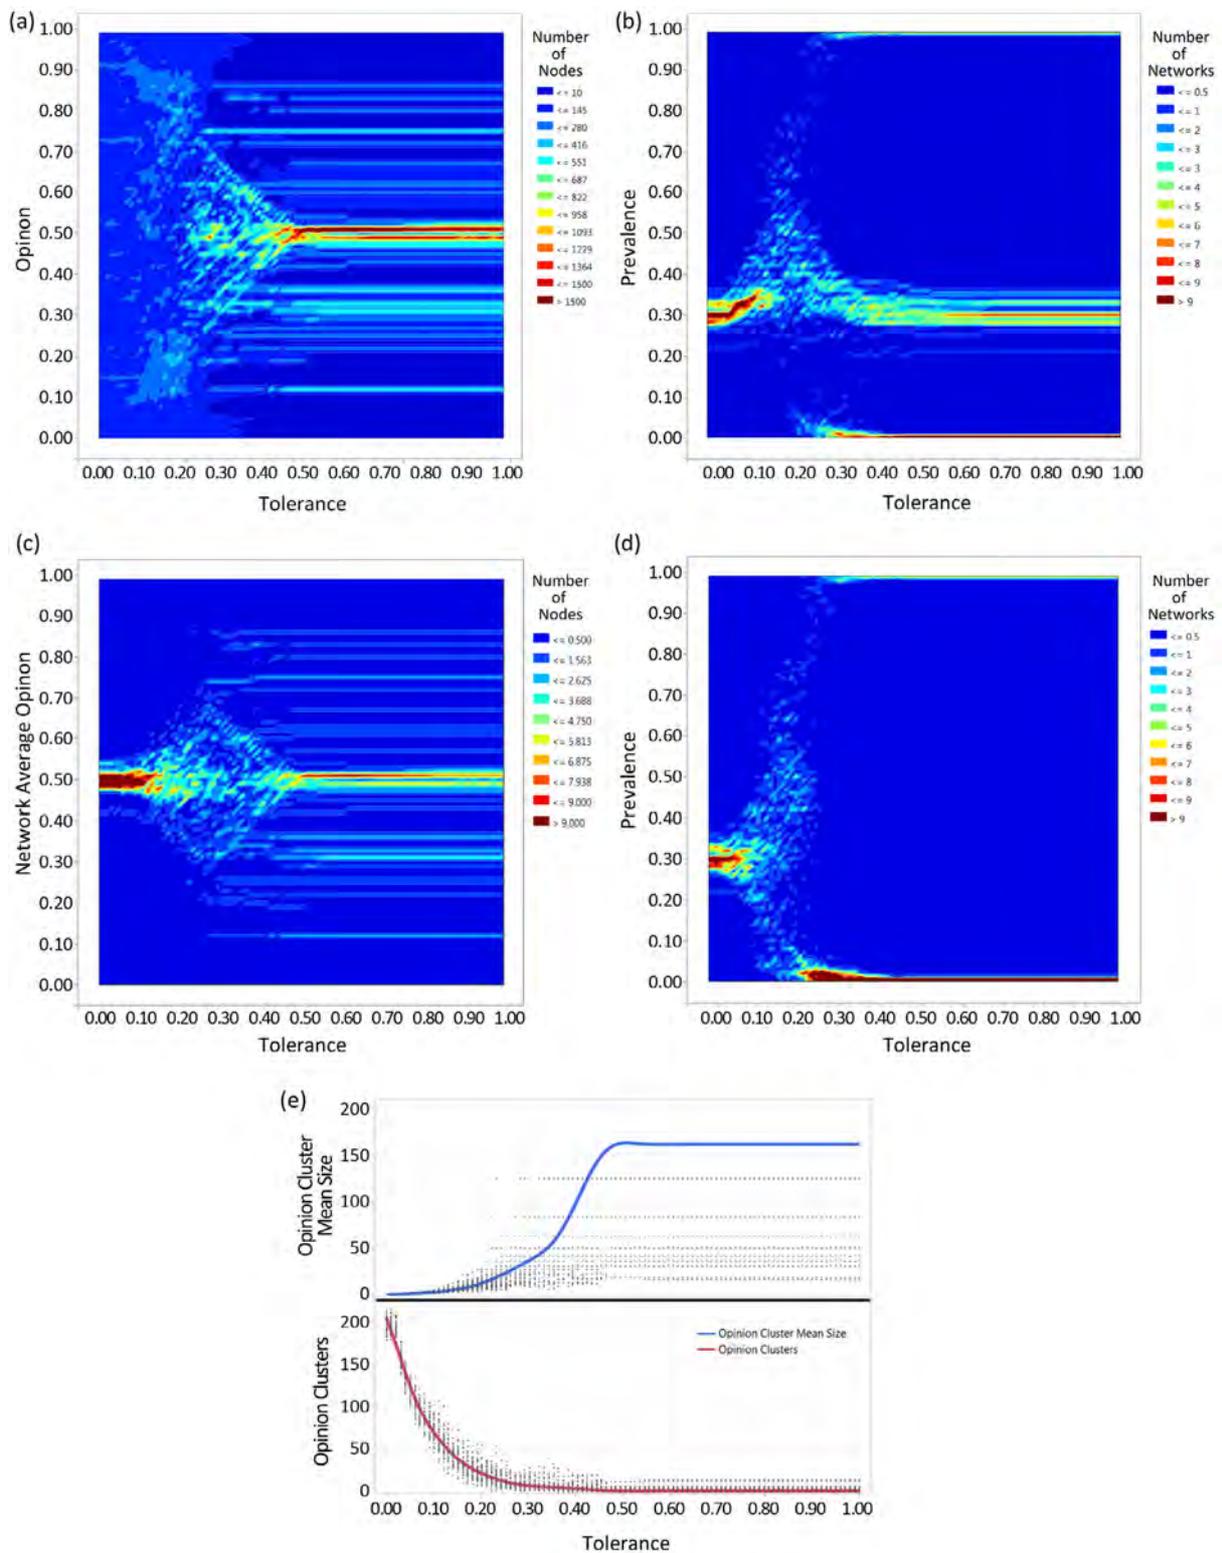


Fig. 12. Opinion (a and c), prevalence (b and d) and opinion cluster sizes and numbers (e) as a function of the tolerance parameter. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run; (c) is a contour diagram showing the average final opinion for each network in each model run. (d) is similar to (b) except that the hysteresis loop (the addiction factor) has been removed. Data from 10 000 model runs are used for each figure.

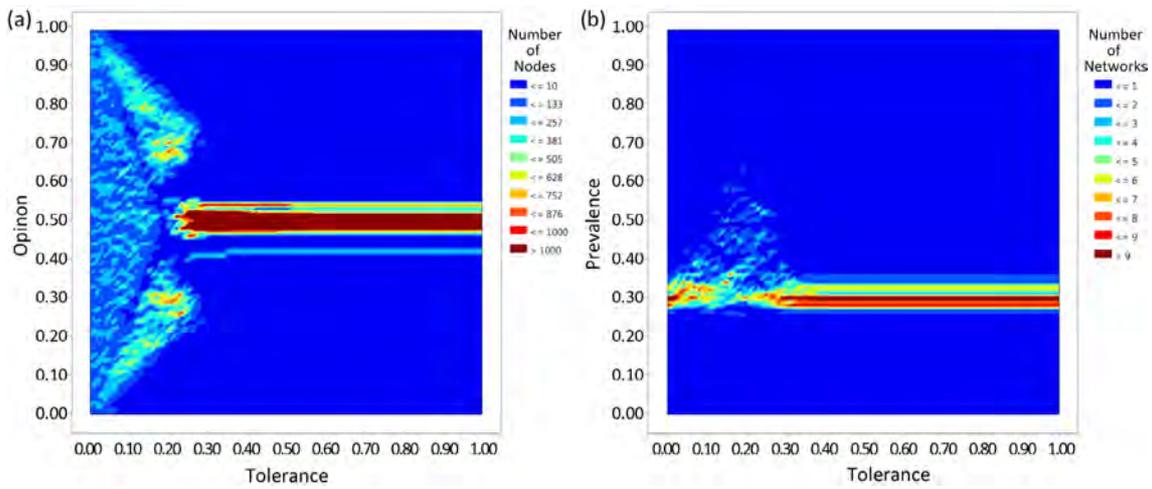


Fig. 13. Opinion (a), and prevalence (b) as a function of the tolerance parameter at a network density of 0.10. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. Data from 10 000 model runs are used for each figure.

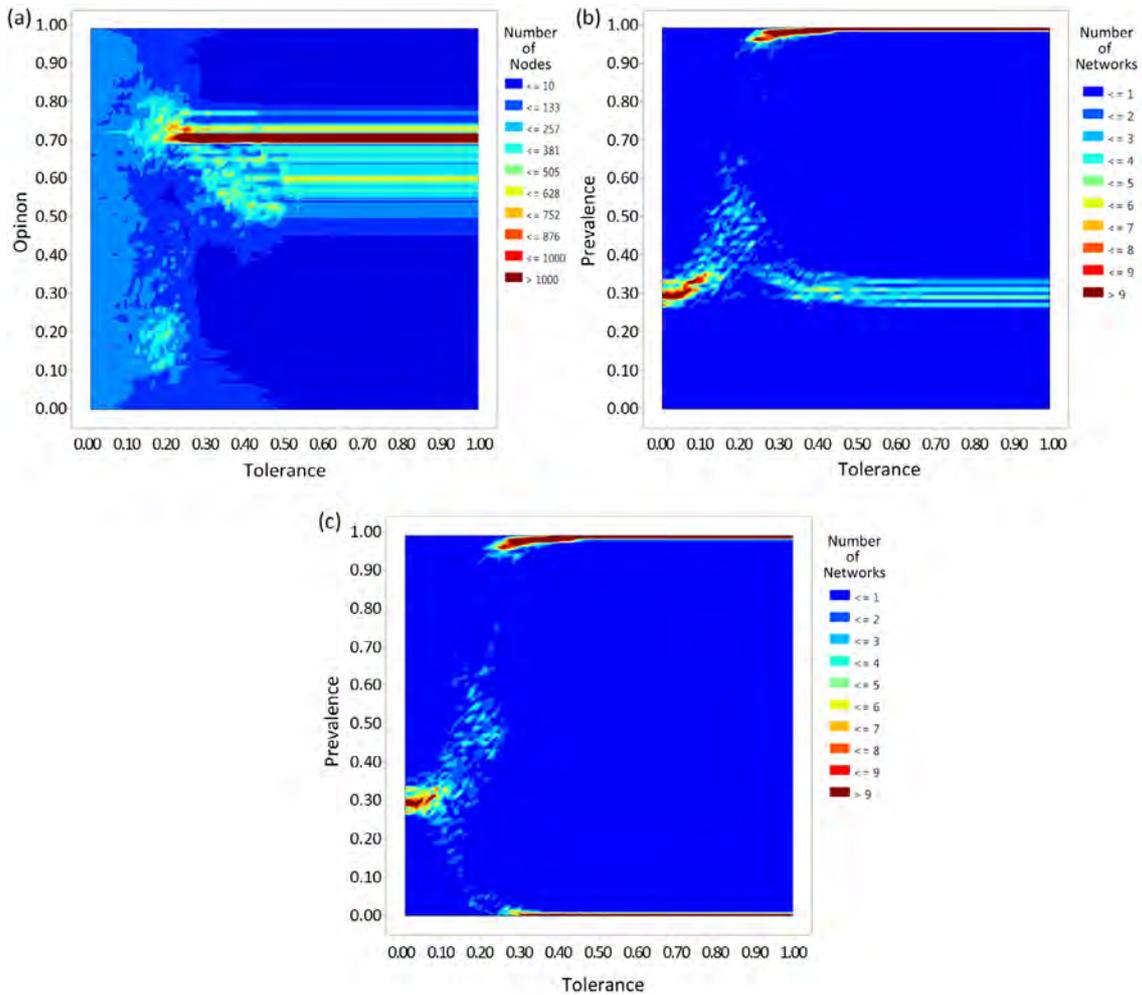


Fig. 14. Opinion (a), prevalence (b, c) as a function of the tolerance parameter in the presence of a media node broadcasting an opinion of 0.725. (a) is a contour plot representing the numbers of nodes holding a given opinion value; (b) is a contour diagram showing the final prevalence for each network in each model run. (c) is similar to (b) except that the hysteresis loop (the addiction factor) has been removed. Data from 10 000 model runs are used for each figure.

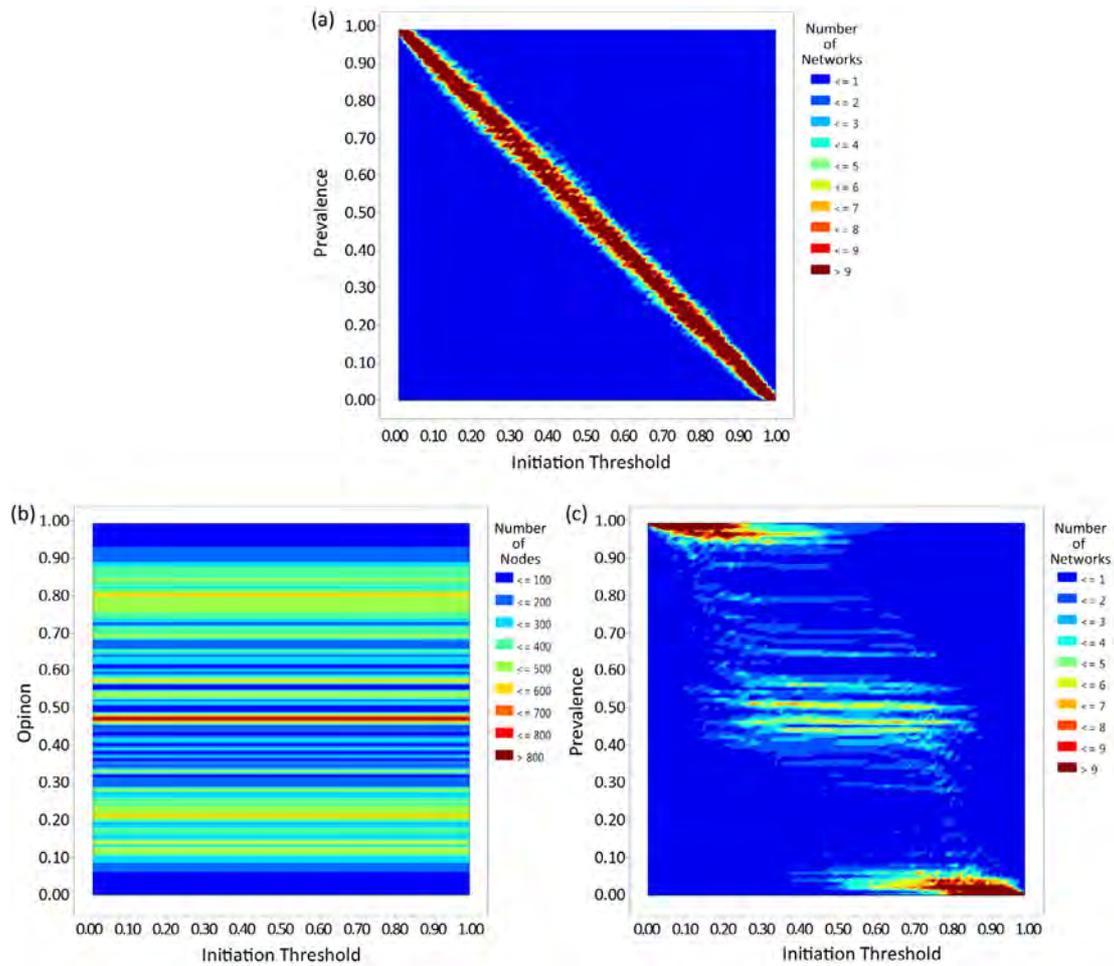


Fig. 15. Prevalence (a) for networks with uniform opinion distributions and no opinion interactions; opinion (b) and prevalence (c) for networks with opinion dynamics and no addiction. All are shown as a function of initiation threshold. Data from 10 000 model runs are used for each figure.

threshold. Increases in the initiation threshold might reflect an increase in monetary price (by raising taxes or setting minimum package sizes), an increase in enforcement of youth access laws, or implementation of clean indoor air policies. Over 80% of current smokers initiated smoking before age 20, with half of those initiating before age 16 [18], so interventions to prevent smoking initiation are primarily targeted to youth and young adults.

The following scenarios illustrate the effects of an increasing initiation threshold on smoking prevalence simulated on 100 randomly generated networks of 250 nodes using 10 000 runs. For some runs one or two media nodes were implemented. A public education node opinion, if present, was set to 0.025 below the effective cessation threshold induced by the addiction factor, and an advertising node opinion, if present, was set to a value 0.025 above the initiation threshold. As indicated above (Sect. 3.2.4), broadcast opinions slightly above the initiation threshold, or below the cessation threshold, are the most effective in eliciting behavioral change. These nodes were attached to networks generated using the default values given in Table 1.

3.2.8.1 Effect of initiation threshold with media nodes

The opinion dynamics algorithm which induces bounded consensus on connected nodes and the addition of addiction-induced hysteresis to the opinion-behavior mapping function introduces more complex behaviors to the otherwise expected linear relationship between initiation threshold and prevalence. Without the addition of opinion dynamics (that is, on a network where tolerance is set to 0), individuals would remain within a uniformly random distribution of opinions, and a 10% increase in the initiation threshold would have a corresponding linear decrease of 10% in initiation, as shown in Figure 15a. Without addiction the decrease of an individual's opinion below the initiation threshold would induce quitting, but the systemic response may still exhibit non-linearities depending on the distribution of opinions in the network. For example, a bimodal network with half the population holding opinions of 0.3 and the other half holding opinions at 0.8 would exhibit no response to an increase in the initiation threshold from 0.6 to 0.7, but would experience a rapid decrease in prevalence as the threshold is increased through 0.8, so long as addiction is not a factor. Figure 15c

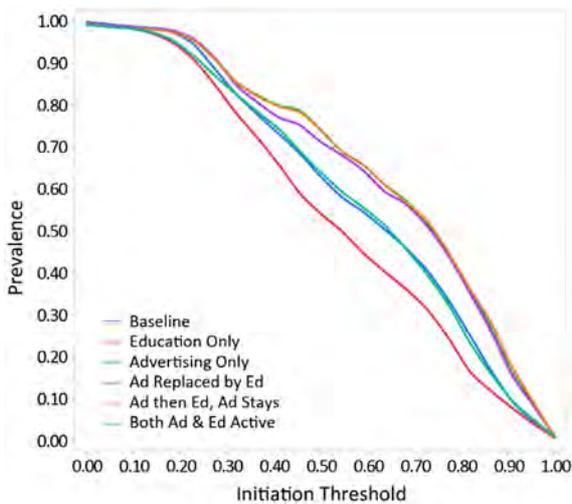


Fig. 16. Prevalence versus initiation threshold for multiple media node scenarios.

shows this response for 0 hysteresis on networks generated with default parameters. The node opinion contour map (Fig. 15b) shows minor peaks at 0.2 and 0.8 (due to the tolerance value of 0.21) and a major peak at 0.5. This produces the pattern of high, moderate, and low prevalence seen in the prevalence contour map (Fig. 15c).

Media nodes affect initiation by attempting to move opinions above or below the behavior threshold. At either extreme of the initiation threshold range, media nodes play only a small role at best in changing prevalence due to the difficulty in attaining network-wide extremal values, as demonstrated in Section 3.2.4. It is in the range of moderate initiation values, from approximately 0.3 to 0.8, that the major differential effects between interventions are observed, as illustrated in Figure 16. This is due to the tendency for network opinions to consolidate near this mid-range.

Figure 16 shows the differential impact of initiation thresholds on prevalence under multiple scenarios, given a hysteresis factor of 0.3. Scenario results cluster into 3 groups. Lowest prevalences are achieved when an educational media node alone is implemented. Somewhat higher prevalences are shown for both the baseline scenario and for a scenario in which advertising and educational media nodes are both active simultaneously. Higher prevalence is reached when an advertising node is connected, the network reaches equilibrium, and then the advertising node is replaced with an educational node. Only slightly higher prevalences are achieved when the advertising node remains active in the presence of the educational node, or when the educational node is not implemented at all. These dynamics are due to the interacting effects of tolerance and the formation of opinion modes under the influence of the advertising node. The three curves at the highest prevalences indicate that if advertising precedes other actions and the network stabilizes with the advertisement, then subsequent educational nodes are far less effective than if they had been implemented at the start of the simulation. The advertising node is more effective due

to its early exclusive presence, which allowed it to drive the network to steady state uncontested. Each scenario is discussed in detail below.

3.2.8.2 Baseline scenario

The baseline scenario for examining the effects of the initiation threshold is shown in Figure 17. The node opinion contour map (Fig. 17a) is highly similar to the one shown above (Fig. 15b), but the default hysteresis factor of 0.3 introduces a change in behavior space. Without hysteresis, behavior (prevalence) exhibits three modes, corresponding to the low, moderate and high opinion modes (Fig. 15c). With hysteresis, individual behavior is not as responsive to changes in initiation threshold or network influences (Fig. 17b). This increases the number of networks with higher prevalence for a given initiation threshold compared to the addiction-free scenario. Hysteresis effectively dampens the network effects, causing a quasi-linear response to changes in the initiation threshold.

3.2.8.3 Educational node scenario

Figure 18 illustrates the effects of the initiation threshold under a scenario in which only an educational node is active. The difference between the initiation threshold and the hysteresis factor induces an effective cessation threshold in the model. For comparability between scenarios, we continue to hold the addiction factor constant at 0.3, and set the education node broadcast opinion to 0.025 below the cessation threshold. This value was chosen based on the media node opinion sweep scenario, investigated above (Sect. 3.2.4), in which broadcast values slightly above or below threshold values achieved greatest behavioral change. In these scenarios, when the cessation threshold is less than 0, the educational node opinion is set to a constant 0, resulting in constant node opinion space for initiation threshold values less than 0.325 (Fig. 18a). After this point, the educational node broadcast opinion increases in step with the increasing cessation threshold, resulting in a linear shift in opinion space. Changes in prevalence are similar in form to the baseline scenario; however, the width of the distributions differ due to the focusing effects of the media node, and average prevalence at each value of initiation threshold is lower for the educational node scenario as shown in Figure 16.

3.2.8.4 Advertising node scenario

Figure 19 shows the results of a scenario in which only an advertising node was present, with a broadcast opinion set to a value 0.025 above the initiation threshold. The effects on opinion (Fig. 19a) are similar to those shown in the investigation of the media node opinion sweep above (Sect. 3.2.4), in which media nodes were shown to influence individuals with similar opinions, but to have no ability to influence those with extremely different opinions. The ability to influence nodes of low opinion when the initiation threshold is low causes saturation at a prevalence of 1.0 for a wider array of initiation threshold values

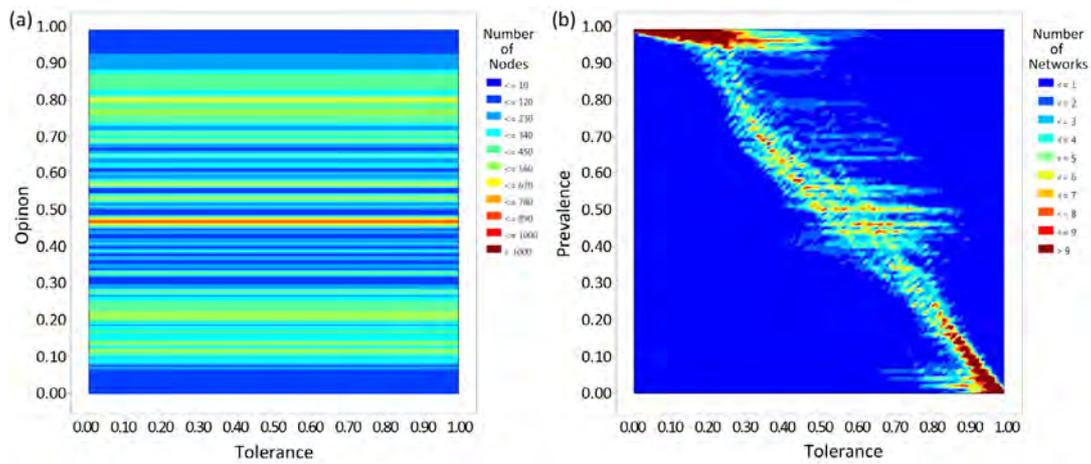


Fig. 17. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in the baseline scenario. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

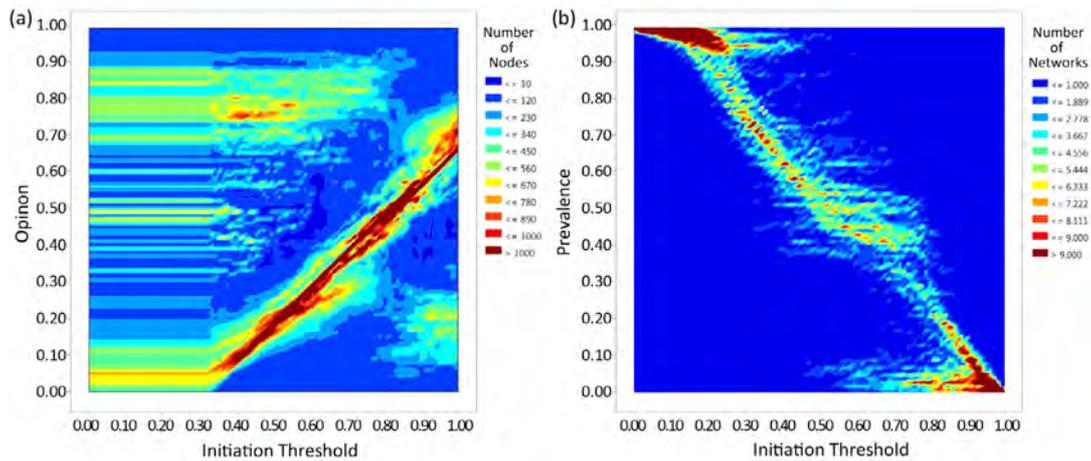


Fig. 18. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in the educational node scenario. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

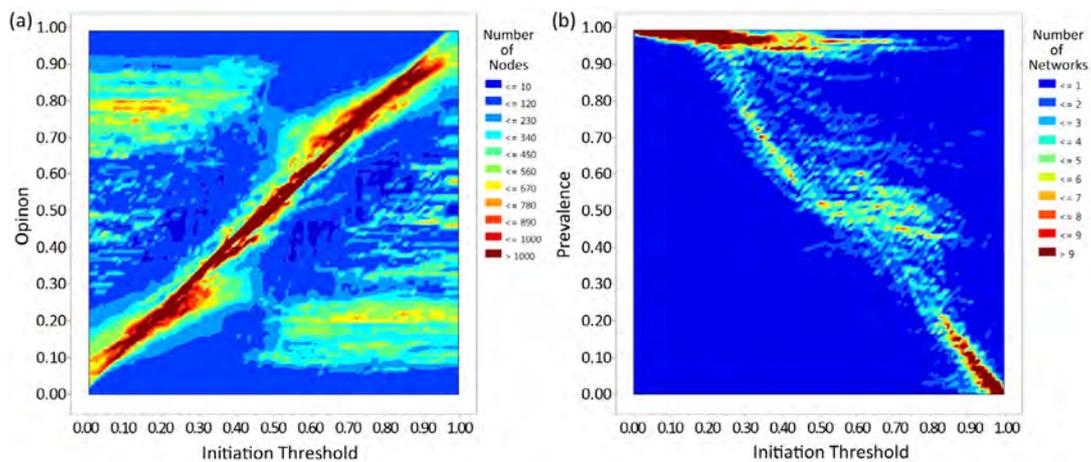


Fig. 19. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in a scenario in which an advertising node is tied to the initiation threshold. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

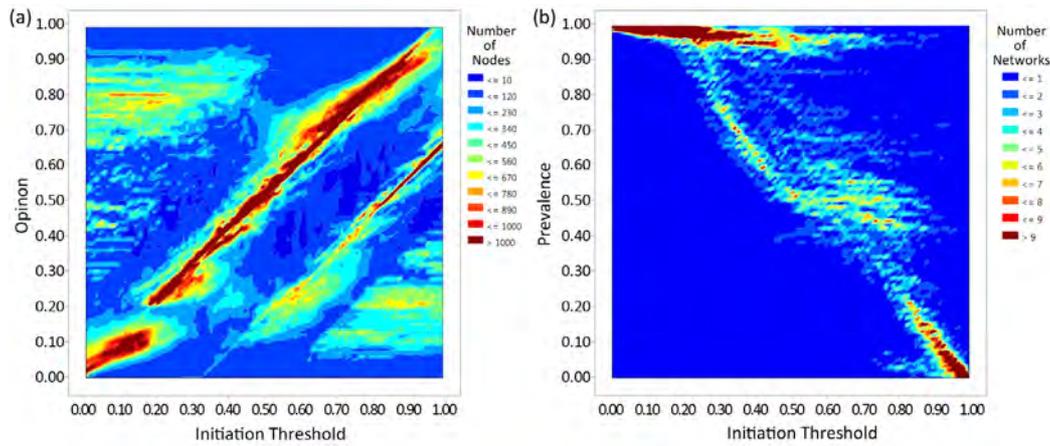


Fig. 20. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The advertising node is initially active, then the educational node also becomes active. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

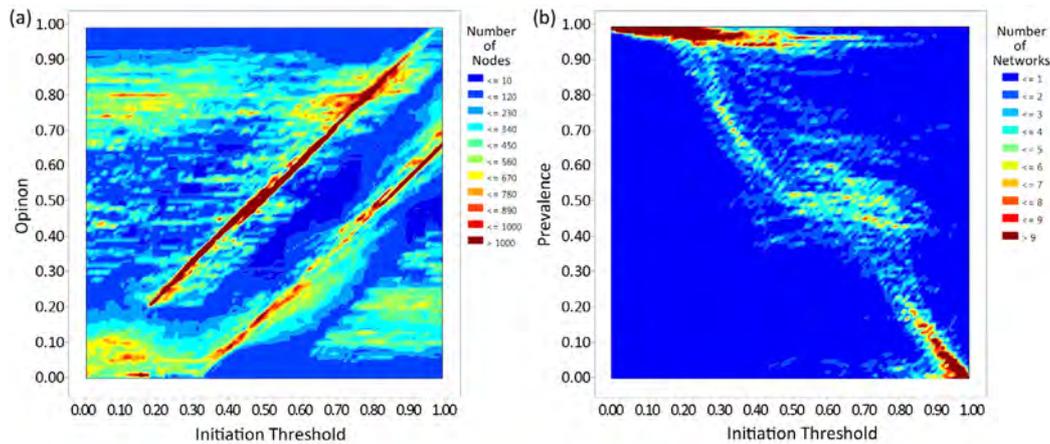


Fig. 21. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The advertising node is initially active, then is replaced with an educational node. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

than was shown in previous scenarios, and overall levels of prevalence are increased (Fig. 19b).

3.2.8.5 Advertising then add education scenario

In Figure 20, we show the results of a scenario in which advertising is active until the network reaches steady state, after which an education media node is introduced. Advertising and education broadcast opinion values are configured as previously, at 0.025 above and below the initiation and cessation thresholds, respectively. In the opinion space contour map (Fig. 20a), we observe the distinct effects of the advertising and education nodes. The stronger band represents the advertising node and follows the trend shown in Figure 19. The lesser diagonal band (Fig. 20a) follows the trend for the educational node shown in Figure 18. The advertising node shows greater effectiveness due to its early exclusive presence, which allowed it to drive the network to steady state uncontested. This results

in a prevalence response extremely similar to that seen in the previous scenario of advertising only (Fig. 19b), indicating that the education media node is less effective if there has been previously unopposed advertising. We again note that for real-world scenarios the situation is considerably more complex, and changes in the network and influence structure would occur that are not represented here.

3.2.8.6 Advertising then replace with education scenario

In Figure 21, we show the results of an identical scenario, except this time removing the advertising node after the network has been driven to steady state under its influence. The resulting opinion space contour map resembles that of the previous scenario, but with the difference that the now uncontested education node has a slightly greater effect. Nevertheless, the advantage of having driven the network to its initial steady state results in

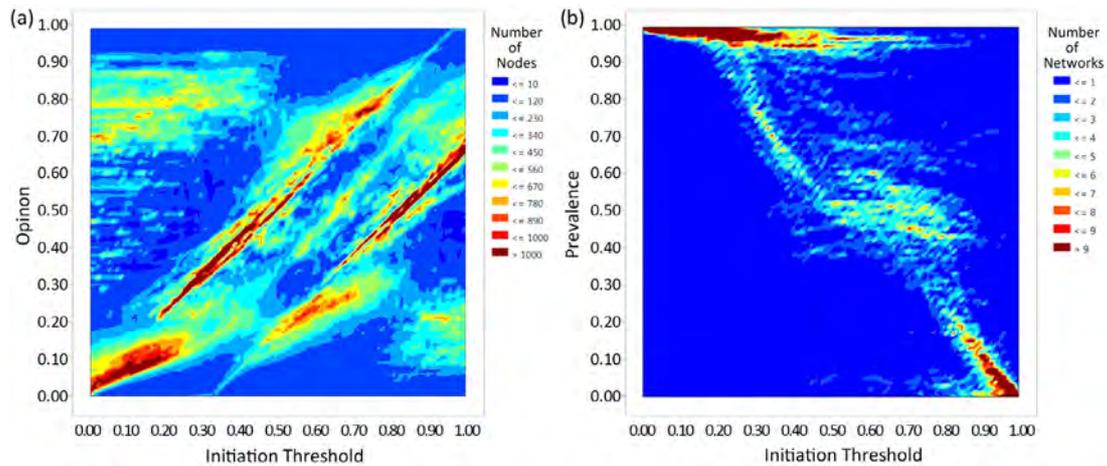


Fig. 22. Opinion (a) and prevalence (b) as a function of initiation threshold for networks in a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The advertising educational nodes are both active throughout. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

the effects of the advertising node on opinion remaining distinctly stronger. This results in a behavioral space that again is nearly identical to that achieved in the advertising only scenario. Tolerance effects constrain the effectiveness of the education node, serving to maintain a collection of high opinion nodes at every point of the run, despite removing the advertising broadcast node. In real-world scenarios, changes in the network structure would occur over time and different individuals might be exposed to different messages. Nevertheless, the results show that providing information early results in an advantage in swaying opinion.

3.2.8.7 Simultaneous advertising and education scenario

In Figure 22, we show the results of a scenario in which both advertising and education are simultaneously active from the start. In this scenario, the opinion space contour map (Fig. 22a) shows greater equality between the modes induced by the advertising and educational media nodes. In addition, a third, less well-expressed mode is introduced between the education and advertising opinions. This results in a behavior space with less consolidation than any of the previous scenarios.

3.2.9 Addiction effect interventions

For the smoking example, the difference between the initiation threshold and the addiction factor, which establishes an implicit cessation threshold for an individual agent, represents the hysteresis in opinion-to-behavior mapping brought about by addiction. This hysteresis serves to maintain smoking behavior despite the fact that an individual's opinion about smoking has fallen. These individuals are effectively latched into a behavior and, depending on the level of physiological and psychological addiction, need to have a far lower opinion of smoking in order to be able to quit than would be necessary without addiction.

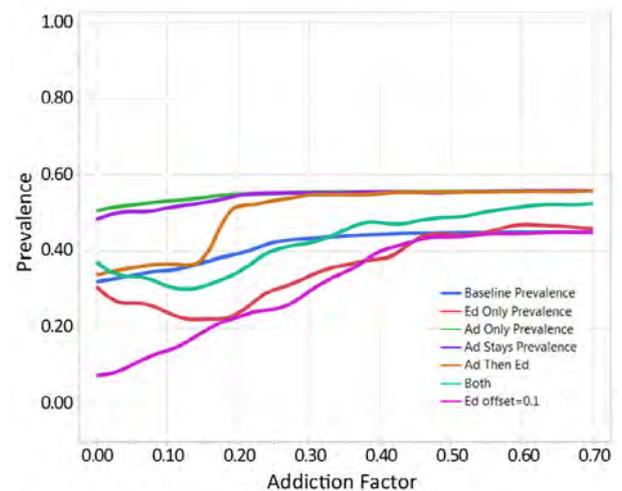


Fig. 23. Prevalence versus addiction threshold for multiple media node scenarios.

Because the opinion-to-behavior function hysteresis establishes a cessation threshold as an offset of the initiation threshold, cessation can similarly be affected by comprehensive tobacco control interventions, such as price increases and smoke-free policies that increase the costs (both monetary and non-monetary) associated with smoking. However, the magnitude of the hysteresis itself can be reduced by interventions that target the effects of addiction, such as providing reduced costs for nicotine replacement therapies and other treatments, as well as social support programs such as quit lines.

3.2.9.1 Comparison of scenarios

For ease of comparison, these results are summarized in Figure 23, which shows the mean prevalences achieved under each scenario. More detailed plots and analysis follow

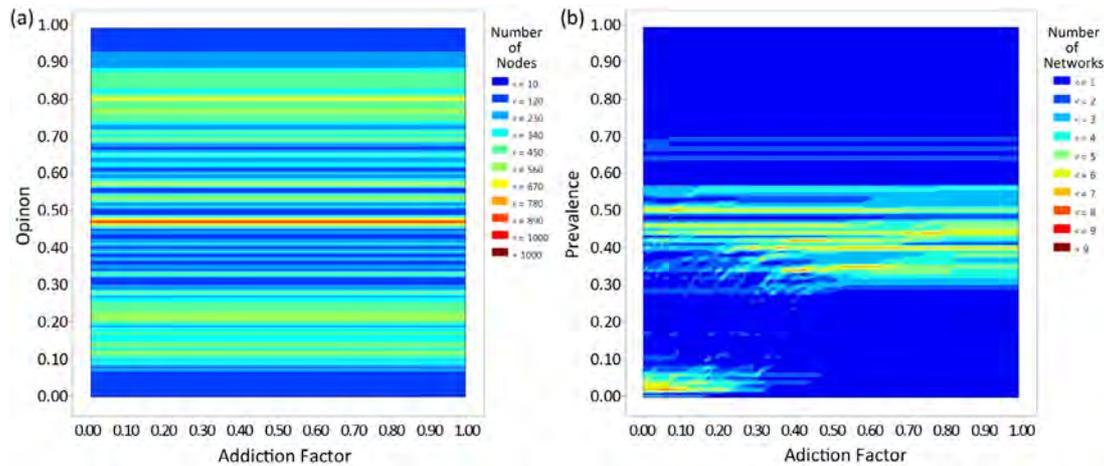


Fig. 24. Opinion (a) and prevalence (b) as a function of the addiction factor for the baseline scenario. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

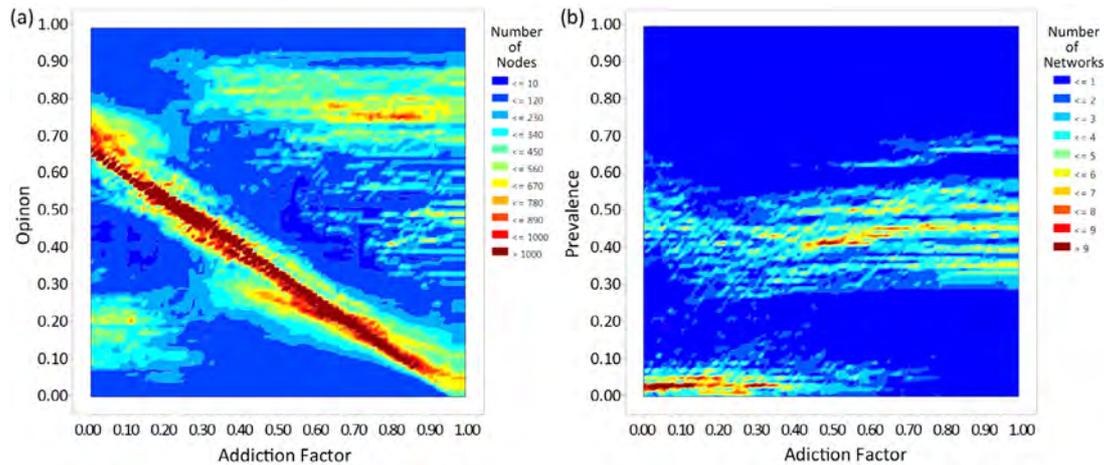


Fig. 25. Opinion (a) and prevalence (b) as a function of the addiction factor for a scenario in which an educational node is tied to the addiction factor. The educational node is present throughout the scenario. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

in Section 3.2.9.3 and later sections. In the baseline scenario, we observe that the effects of increasing the addiction factor largely saturate above 0.25. This is due to the consolidation of the major mode with opinion values around 0.5, shown in Figure 24a. Individuals in the central mode are either never smokers (their opinions were never above the initiation threshold), or else their opinions were lowered from above 0.7. If their opinions were lowered from above the initiation threshold, and if the addiction factor is larger than 0.25, then they will likely be caught in hysteresis, thus raising prevalence. Although another major mode exists in Figure 24a below an opinion value of 0.3, these individuals are generally never smokers, and so do not have their behaviors influenced by addiction.

At very low addiction factor values, the education node’s broadcast opinion is close to the initiation threshold. For example, an addiction factor of 0 (no addiction) sets the effective cessation threshold to 0.7 (equal to the

initiation threshold), and the education node broadcast opinion is correspondingly set to 0.675. This value is not sufficiently below the initiation threshold to have a large effect on prevalence, and so the prevalence is the same as seen in the baseline scenario. As the addiction factor increases and the education node broadcast opinion is correspondingly lowered, effectiveness in reducing prevalence increases until the broadcast opinion is low enough (at an addiction factor of 0.2, corresponding to a broadcast opinion value of 0.475) that an uninfluenced, higher opinion mode begins to form, as shown in Figure 25a. At this point, although the education node scenario achieves a lower prevalence than the other scenarios, the increasing addiction factor results in increasing prevalence. Finally, by the time the addiction factor reaches 0.45 and the education node’s opinion is a corresponding 0.225, it is exiting the range of influence of even mid-opinion smokers, and so does not achieve an outcome significantly different

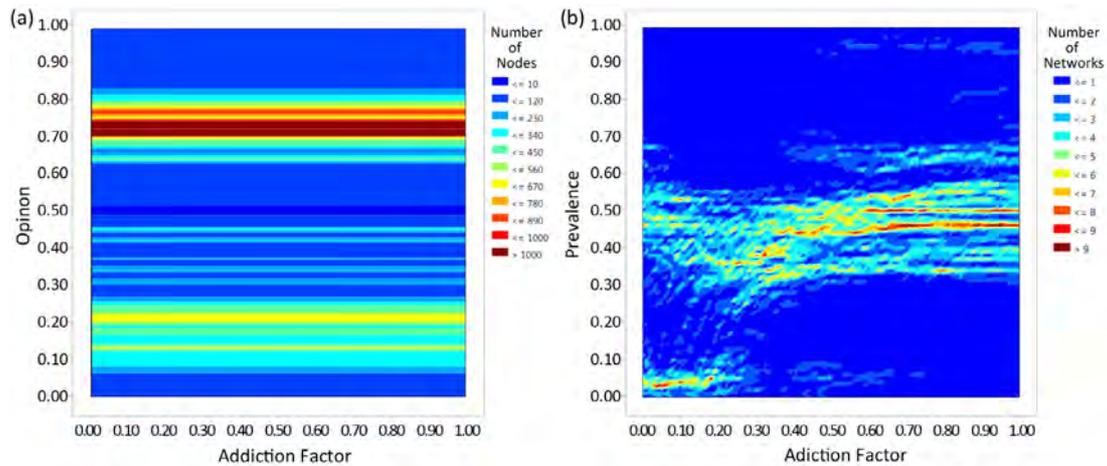


Fig. 26. Opinion (a) and prevalence (b) as a function of the addiction factor for a scenario in which an advertising node is tied to the initiation threshold. The advertising node is present throughout the scenario and does not change value. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

from the baseline scenario. We illustrate that the effect shown in the scenarios containing an education node on the interval $[0.0, 0.2]$ is due to a less effective education campaign, rather than due to the addiction factor itself. When the education media node's broadcast opinion value is offset by 0.1 rather than 0.025, as shown in Figure 23, the education node is effective at reducing prevalence for addiction factors below 0.2. Above this addiction factor, prevalence becomes essentially equivalent to the scenario in which the education node broadcasts an opinion with an offset of 0.025.

In the simultaneous education and advertising scenario, we see a prevalence curve nearly identical in shape to education only, but shifted up by 0.05. Advertising is able to increase prevalence even in the presence of education, but the overall response curve is that of the education node. This is due to the education node's dynamic value being set based on the current level of addiction, while the advertising node's broadcast opinion value is held constant based on the constant initiation threshold. A dynamic educational node with opinion that can change with addiction level is more effective than a constant advertising opinion based on a constant initiation threshold.

Removing the advertising node and adding an education node achieves near-baseline prevalence values only for low levels of addiction. The stronger the effects of addiction, the less effective it becomes to replace advertising with education. At an addiction factor of 0.15 and an effective cessation threshold of 0.55, the education node's opinion becomes 0.525, which is outside the opinion range of the higher opinion nodes, allowing the formation of a high opinion mode between 0.7 and 0.85, as seen in Figure 27. This causes a phase transition in the prevalence curve, yielding prevalence equivalent to that seen in the other scenarios incorporating advertising due to the effects of tolerance; at higher addiction levels the educational node moves outside of the tolerance range of more individuals.

For the scenarios in which advertising remains active throughout a run (Advertising only and Ad Stays Active), there is little difference with or without a late-introduced education node. As shown in Figures 26a and 29a, opinion space is sufficiently biased by the early-acting advertising node that even though the education node shifts some opinions at low addiction factors, opinion shifts are small due to tolerance bound constraints and effects on prevalence are minimal. In real-world dynamic networks, in which births, aging and deaths modify membership and constantly changing messaging campaigns introduce new sources of opinions, message campaign results can vary. For example, Hamilton [40] asserts that the Fairness Doctrine, which required broadcasters to balance cigarette advertising with educational messages on smoking-induced harm, caused a decrease in prevalence despite long-term industry advertising.

3.2.9.2 Impact of addiction and cessation threshold

The effects of sweeping through values for the addiction factor are strikingly different than those of sweeping through initiation thresholds. While the effects of addiction are relatively simple for an individual node, the system as a whole displays complex behavior whose origin becomes clear when considering the interaction between opinion and behavior as moderated by addiction and driven by media nodes.

3.2.9.3 Baseline scenario

Figure 24 shows the effects of a sweep of addiction factors for the baseline case. With no media nodes driving network opinions, opinions remain constant with increasing addiction, with a highly-expressed mode at opinions around 0.5 and lower-expressed modes around 0.2 and 0.8 (Fig. 24a). Given the baseline initiation threshold of 0.7,

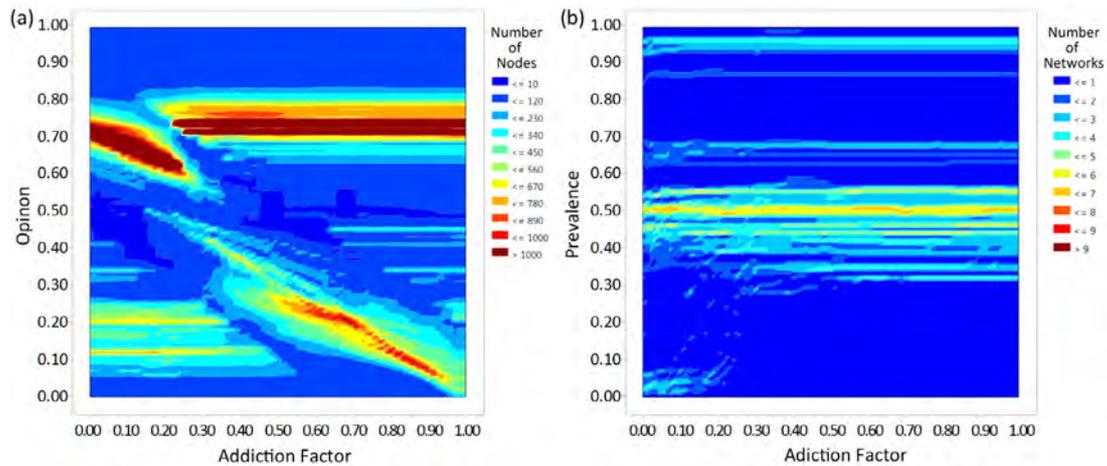


Fig. 27. Opinion (a) and prevalence (b) as a function of the addiction factor for a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The advertising node is present throughout the scenario and the educational node is introduced after the networks reach equilibrium. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

addiction factors below 0.2 allow for cessation by smokers with opinions of 0.5 or lower, as reflected by the low-prevalence cluster in Figure 24b. If the addiction factor increases above 0.2, smokers need an opinion below 0.5 in order to quit. However, the smokers in the network are primarily made up of individuals whose opinions at steady state are above the initiation threshold, or those whose opinions were formerly above the initiation threshold. Because opinion values in this scenario trend towards 0.5 for both initial smokers and non-smokers, the lower opinions achieved by formerly-high-opinion nodes are insufficient to allow cessation, thus contributing to the relatively unchanging prevalence values above approximately 0.2.

3.2.9.4 Educational node scenario

Figure 25 shows the effects of a sweeping on the magnitude of hysteresis (the addiction factor) under the influence of an education node broadcasting an opinion value 0.025 below the cessation threshold. This value was selected based on results shown in Section 3.2.4. At an addiction factor of 0, the effective cessation threshold is equal to that of the initiation threshold, which in this case is 0.7. The education node broadcast opinion is therefore set to 0.675. At this value, the education node is successful at influencing high-opinion individuals, but it does so by broadcasting an opinion value very close to the initiation threshold, which lowers its overall effect on reducing prevalence. As the education node broadcast opinion decreases with increasing addiction and gains separation from the initiation threshold, its effectiveness increases. However, after the addiction factor reaches 0.2, the education node's broadcast opinion of 0.475 becomes too low to effectively influence the opinions of the high-opinion nodes, again causing a loss in effectiveness and higher prevalence (Fig. 25b). Continuing to lower the education node's broadcast opinion results in increasing alienation of high- and even moderate-

(opinion) nodes throughout the rest of the scenario. However, these extremely low media node opinion values are necessitated by the effective cessation threshold induced by the hysteresis factor – higher broadcast opinion values would place the education node's opinion above that of the cessation threshold, and thus would potentially influence opinions but not enough to trigger a behavioral change. Thus, as the education node moves to lower opinion values, it stops influencing people with high opinions. The increase in prevalence is due to both the increase in addiction and the effects of tolerance-induced constraints on the education node's ability to change opinions.

3.2.9.5 Advertising node scenario

Figure 26 shows the results of addiction under a scenario with only an advertising node. Because advertising node opinion is fixed at 0.025 above the initiation threshold of 0.7, and because the initiation threshold is held constant in these scenarios, there is no change in node opinions as a result of increasing the addiction factor (Fig. 26a). We observe that the advertising node induces the formation of a highly expressed mode around 0.725, and a more modestly expressed lower mode around 0.17 comprised of nodes that are unaffected by the advertisement due to tolerance constraints. This concentration of high opinion nodes reduces the effectiveness of low addiction factors in reducing prevalence (Fig. 26b), rendering them all but ineffectual even at the lowest levels.

3.2.9.6 Advertising then add education scenario

Figure 27 shows the effects of giving the advertising node a head start – allowing the network to come to steady state under the effects of advertising before introducing the dynamic opinion education node described in the above education-only scenario. Contrasted with that scenario,

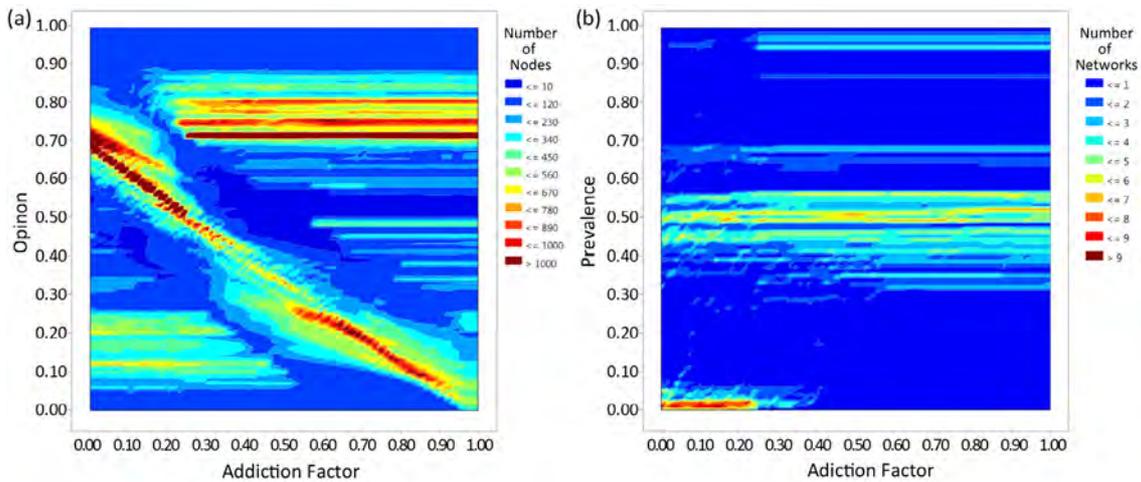


Fig. 28. Opinion (a) and prevalence (b) as a function of the addiction factor for a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The advertising node is present until each network reaches equilibrium, then it is replaced with an educational node. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

the upper mode is far more highly expressed, due to the effects of the advertising node (Fig. 27a). At high education node broadcast opinion values, used when the addiction factor is very low, the education node is effective at influencing high-opinion nodes, and is able to lower those opinions as its opinion decreases and addiction increases. However, because the education node opinion is close to the initiation threshold and therefore similar to that of the advertising node, it has only a modest effect on prevalence, showing only a slight decrease over that of the advertising-only scenario (Fig. 27b). Above an addiction factor of 0.2, corresponding to an education node opinion of 0.525, even this modest effect is lost as the education broadcast opinion passes outside of the tolerance threshold of the high-opinion nodes. A larger tolerance value would permit a greater range of influence for the education node.

3.2.9.7 Advertise then replace with education scenario

Figure 28 shows the results of a scenario in which the advertising node is removed after the initial steady state is reached, giving the education node an advertising-induced steady state as its initial condition in an otherwise uncontested environment. Contrasted with the previous scenario in which the advertising node remains active, we observe that the advertising-induced upper opinion mode, while more strongly expressed than in the education-only scenario, is far less highly expressed than when the advertising node remains active (Fig. 28a). This scenario allows the education node greater ability to influence opinions, with a resulting improvement in reducing smoking behaviors (Fig. 28b). However, this benefit is lost at lower values of education broadcast opinion that are outside the tolerance range of the upper opinion nodes; there, the network remains closer to the prevalence levels seen in the advertising-only scenario.

3.2.9.8 Simultaneous advertising and education scenario

Finally, Figure 29 illustrates the effects of an even-handed scenario in which both advertising and education nodes are broadcasting at the same time. The contour map of node opinion (Fig. 29a) shows characteristics of both the previous scenario and the advertising-only scenario, but the ability of the education node to influence opinions is increased because it begins its activity on a network that has not come to steady state under the influence of advertising. This results in a behavior-space reaction (Fig. 29b) similar to the education-only scenario, with nearly identical transitions but with slightly reduced levels of effectiveness.

4 Conclusions

Opinion dynamics models are powerful tools for modeling complex phenomena in the social sciences. With the addition of behavioral components, the scope of applications for this category of models expands significantly. Such applications can be used to analyze the fundamental dynamics of a system as well as to model effects of a variety of interventions under a range of counterfactual scenarios.

This paper focused on the effects of idealized network topological elements on opinion dynamics and hysteresis in behavior and how they impact the potential distribution of opinions and network behavior states for a given set of parameter values. Opinion dynamics models using well-mixed populations with random connections and fully-connected networks converge toward single opinion values for tolerances greater than 0.27. With sparse connectivity the network remains in non-consensus over a broader range of tolerance values. In the case of the ER-type networks with probability of node connection of 0.022, consensus is not reached for tolerances less than 0.5. When the

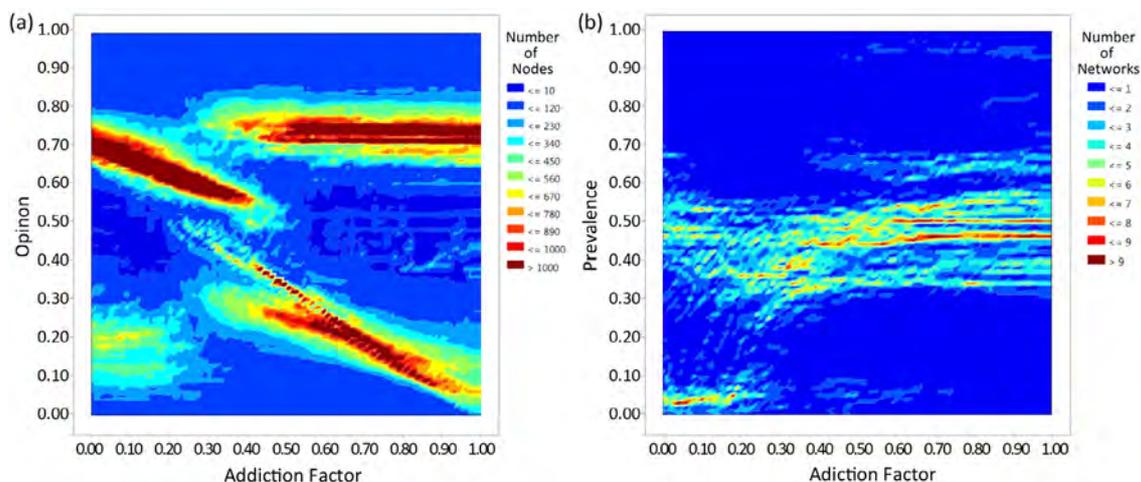


Fig. 29. Opinion (a) and prevalence (b) as a function of the addiction factor for a scenario in which an advertising node is tied to the initiation threshold and an educational node is tied to the addiction factor. The both nodes are active throughout the scenario. (a) is a contour diagram showing the final opinion for each node in each model run; (b) shows the distribution of prevalences. Data from 10 000 model runs are used for each figure.

probability of node connection is increased to 0.2, the outcomes are similar to previous studies with a tolerance threshold for consensus of 0.25; but with an added phase of trimodal results for tolerance between 0.175 and 0.25.

Mapping opinion to a discrete, bimodal behavior results in fewer modes for the node behavior state than for the opinion state. Aggregating the behavior state of the nodes into a network prevalence value further reduces the modality of the model results. The behavior of the model is strongly affected by initiation and cessation threshold values. Increasing the initiation threshold, which, for our example, can be interpreted as raising the costs associated with starting to smoke, often results in a roughly linear decrease in the number of new smokers, although the shape of the response curve is driven by the distribution of opinion in the network.

Hysteresis in the opinion-to-behavior mapping dampens network and intervention effects on individual behavior. With regard to our smoking example, this suggests that it will be easier to prevent individuals from starting than to get people to quit, consistent with evidence showing that while the majority of smokers in the U.S. express an interest in quitting and have tried to quit, only a small fraction are successful [29]. Both sets of effects can be altered by the activity of media campaigns that can raise or lower opinions. These model runs also highlight the potential value of this modeling approach for exploring and evaluating multi-component interventions focused at different levels of influence.

This work benefitted from the programming skills of Ryan Hammer and Taylor Berger and the technical editing and advice of Louise Maffitt. The authors thank each of these individuals for their contributions. The authors also wish to thank Brian Rostron, Shawn Fultz Conrad Choiniere and David Portnoy for on-going insight and support.

This work was funded by the U.S. Food and Drug Administration through a contract with the U.S. Department of Energy/Sandia National Laboratories (funding document 224109011). 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. The funders, other than employees named as authors or described in the acknowledgments, had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. This article represents the views of the authors, and should not be construed to represent FDA's views or policies.

References

1. G. Weisbuch, G. Deffuant, F. Amblard, J.P. Nadal, *Complexity* **7**, 55 (2002)
2. S. Galam, *Physica A* **336**, 49 (2004)
3. A. Martins, *Phys. Rev. E* **78**, 036104 (2008)
4. K. Sznajd-Weron, J. Sznajd, *Int. J. Mod. Phys. C* **11**, 1157 (2000)
5. S. Galam, *Int. J. Mod. Phys. C* **19**, 409 (2008)
6. J. Lorenz, *Int. J. Mod. Phys. C* **18**, 1819 (2007)
7. S. Galam, in *Sociophysics: A Physicist's Modeling of Psycho-Political Phenomena* (Springer, Boston, 2012), pp. 169–202
8. G. Weisbuch, *Eur. Phys. J. B* **38**, 339 (2004)
9. C. Castellano, S. Fortunato, V. Loreto, *Rev. Mod. Phys.* **81**, 591 (2009)
10. D. Cartwright, F. Harary, *Psychol. Rev.* **63**, 277 (1956)
11. R. Hegselmann, U. Krause, *Comput. Econ.* **25**, 381 (2005)
12. J.R.P. French Jr., *Psychol. Rev.* **63**, 181 (1956)
13. A.C.R. Martins, *Int. J. Mod. Phys. C* **19**, 617 (2008)
14. S.E. Asch, *Scientific American* **193**, 31 (1955)
15. N.A. Christakis, J.H. Fowler, *N. Engl. J. Med.* **358**, 2249 (2008)

16. N.A. Christakis, J.H. Fowler, N. Engl. J. Med. **357**, 370 (2007)
 17. J.T. Cacioppo, J.H. Fowler, N.A. Christakis, J. Pers. Soc. Psychol. **97**, 977 (2009)
 18. *Preventing Tobacco Use Among Youth and Young Adults: a Report of the Surgeon General (Three Volume Set)*, edited by P.H.S.U.S, Centers for Disease Control and Prevention (U.S.), Office on Smoking and Health (Public Health Service, 2012)
 19. G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Adv. Complex Syst. **3**, 87 (2000)
 20. R. Hegselmann, U. Krause, J. Artif. Soc. Soc. Simul. **5**, 2 (2002)
 21. C. Alexander, M. Piazza, D. Mekos, T.W. Valente, J. Adolesc. Health **29**, 22 (2001)
 22. National Cancer Institute, *The Role of the Media in Promoting and Reducing Tobacco Use* (US Department of Health and Human Services, National Institutes of Health, National Cancer Institute, Bethesda, 2008)
 23. T.E. Langley, A. McNeill, S. Lewis, L. Szatkowski, C. Quinn, Addiction **107**, 2043 (2012)
 24. T.W. Moore, P.D. Finley, N.S. Brodsky, T. Brown, B. Apelberg, B. Ambrose, R. Glass, J. Artif. Soc. Soc. Simul., to be published
 25. T. Carletti, D. Fanelli, S. Grolli, A. Guarino, Europhys. Lett. **74**, 222 (2006)
 26. Centers for Disease Control and Prevention (CDC), MMWR Morb. Mortal. Wkly. Rep. **59**, 1063 (2010)
 27. C.M. Fichtenberg, S.A. Glantz, British Med. J. **325**, 188 (2002)
 28. C. Evans, J. Fielding, R. Brownson, P. Buffler, M. England, D. Fleming, M. Fullilove, F. Guerra, A. Hinman, G. Isham, G. Land, C. Mahan, P. Mullen, S. Scrimshaw, R. Thompson, R. Lawrence, J. McGinnis, L. Novick, S. Teutsch, Am. J. Prev. Med. **20**, 10 (2001)
 29. Centers for Disease Control and Prevention (CDC), MMWR Morb. Mortal. Wkly. Rep. **60**, 1513 (2011)
 30. G. Fong, D. Hammond, F. Laux, M. Zanna, K.M. Cummings, R. Borland, H. Ross, Nicotine Tob. Res. **6**, 341 (2004)
 31. N.L. Benowitz, N. Engl. J. Med. **362**, 2295 (2010)
 32. Centers for Disease Control and Prevention (US), National Center for Chronic Disease Prevention and Health Promotion (US), Office on Smoking and Health (US), *How Tobacco Smoke Causes Disease: the Biology and Behavioral Basis for Smoking-Attributable Disease: a Report of the Surgeon General* (Centers for Disease Control and Prevention (US), Atlanta, 2010)
 33. United States Public Health Service Office of the Surgeon General, *Reducing the Health Consequences of Smoking: 25 Years of Progress – a Report of the Surgeon General* (US Department of Health and Human Services, Public Health Service, Centers for Disease Control, Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health, 1989)
 34. World Health Organization (2009), http://www.who.int/tobacco/mpower/2009/gtcr_download/en/
 35. L.F. Stead, R. Perera, T. Lancaster, Tobacco Control **16**, i3 (2007)
 36. F.J. Chaloupka, J.A. Tauras, J. Econ. Finan. **28**, 395 (2004)
 37. P. Erdős, A. Rényi, Publ. Math. Debrecen **6**, 290 (1959)
 38. M.E. Newman, S.H. Strogatz, D.J. Watts, Phys. Rev. E **64**, 026118 (2001)
 39. Centers for Disease Control and Prevention (CDC), MMWR Morb. Mortal. Wkly. Rep. **60**, 1207 (2011)
 40. J.L. Hamilton, Rev. Econ. Stat. **54**, 401 (1972)
- Open Access** This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.