TÉCNICO LISBOA

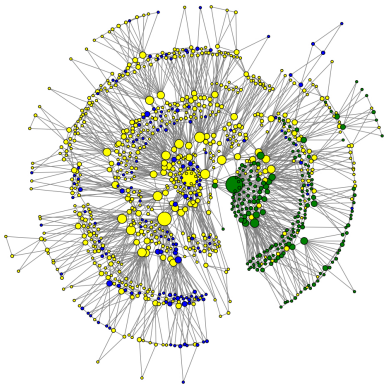

Topological effects in the spread of information and misinformation in social networks

Gonçalo Alexandre Santos Simões

Thesis to obtain the Master of Science Degree in

Computer Science and Engineering

Supervisors: Prof. Francisco João Duarte Cordeiro Correia dos Santos
Prof. Jorge Manuel dos Santos Pacheco

Examination Committee

Chairperson: Prof. Pedro Tiago Gonçalves Monteiro
Supervisor: Prof. Francisco João Duarte Cordeiro Correia dos Santos
Member of the Committee: Prof. Fernando Pedro Pascoal dos Santos

September 2021

Resumo

Notícias falsas são um problema crescente para o mercado de notícias e a desinformação de uma larga parte da população mundial. Este trabalho foca-se no estudo da dinâmica de opiniões e estratégias de intervenção capazes de reduzir a propagação de notícias falsas. Para este fim, desenvolvemos um novo simulador de processos de propagação no qual interações socias são representadas por uma rede de nós conectados. Partindo de três tipos de redes livres de escala, abordámos o impacto de dois tipos de intervenções e dinâmicas de propagação de rumores: 1) método semelhante a vacinação para informar os nós, tornando-os resistentes à informação falsa; 2) competição entre rumores verdadeiros e falsos. Os resultados mostram que em estruturas de baixo agrupamento, a propagação de ambos os tipos de notícias é facilitado. Em estruturas hierárquicas, o fluxo é restringido pelos nós de topo, condicionando a propagação de informação. Finalmente, em comunidades pouco conectadas, a baixa quantidade de conexões entre comunidades causam o atraso do pico de propagação, mas no final são obtidos valores similares aos observados nas estruturas de baixo agrupamento. Além disso, estruturas de comunidades levam a níveis elevados de polarização, dado que criam "bolhas de opinião". Um efeito similar também é possível ser observado em topologias hierárquicas com agrupamento elevado.

Palavras-chave: Propagação de rumores; Intervenção; Notícias Falsas; Redes sem escala;

Abstract

Fake news and misinformation are an increasing problem at local and global scale. This work focuses on the analysis of opinion dynamics and interventions strategies able of reducing the spread of fake news. To this end, we develop a novel simulator of spreading processes where social interactions are described as a network of connected nodes. We resort to different classes of scale-free networks to assess two types of interventions and rumour spreading dynamics: 1) vaccination-like method to inform nodes, making them resistant to fake information; 2) competition between true and fake rumours. We show that in low clustering structures, the spread of both true and fake news is facilitated. In networks portraying hierarchical topologies, the flow of information is largely dominated by top nodes, which suggests the use of targeted intervention policies. Finally, if social networks are shaped by sparsely connected communities, the low inter-community connections lead to a later peak in the rumour spread but still reaches similar final values as in low clustering structures. Furthermore, community structure leads to high levels of polarisation, as a result of the emergence of "belief bubbles". A similar effect may also be observed in hierarchical topologies with high clustering coefficient.

Keywords: Rumour spreading; Intervention; Fake News; Scale-free networks;

Acknowledgements

I want to express my gratitude to Dr. Francisco C. Santos for giving me the opportunity to work on this matter and for helping me throughout the development of this thesis. Without his enthusiasm and encouragement this thesis might never have seen the light of day.

Secondly, I want to thank Dr. Jorge M. Pacheco for his help in finding the right path of research, especially during the beginning when the research topic was still vague, and asking the right questions to help this thesis reach its focus.

I am extremely grateful for my colleagues, and friends, Pedro Duarte and Fábio Vital for their help and collaboration in doing the aforementioned previous work in this area. Pedro Duarte also helped me further during the development and writing of this thesis by helping me make reasonable goals and encouraging me to meet them by promoting regular discussion of the topics/results.

Contents

1 Introduction

Fake news are pieces of information maliciously created or spread without fact-checking. They can also be seen as **rumours** that threaten the public opinion in current matters by creating unnecessary discord [2, 23]. When in excess they can also undermine the credibility of news markets [27].

Nowadays, individuals consume information in a variety of ways: newspapers, websites, aggregators, social networks, etc.; with a clear shift for digital platforms. This pushed a digital transformation of the news markets to adapt to these novel distribution channels. Out of these, social networks are the most used, but they also happen to be the fastest for rumour spreading. While fake news are still a small part of people's news consumption, in general, they are much more prevalent in partisan media – biased and politically aligned media outlets – where they strengthen their own political bias [27]. Taking advantage of this bias and social networks, an increasing number of bots has been used to wage political and opinion wars [8, 23, 33, 36], generating more fake news.

What also defines fake news is how they spread in the network and S. Vosoughi et al. [41] mention that due to the more novel nature of fake news, when compared to true news, people tend to share them more, and they end up reaching farther, faster, deeper and more broadly in the network. They also found that fake news tend to be associated with different emotions than true news, where fake news led to replies with fear, disgust, and surprise, while true news led to anticipation, sadness, joy, and trust. Moreover, their findings also seem to imply that humans tend to share fake news more than bots. Z. Zhao et al. [49] instead focus on the topological features of the spread and find that fake news propagate differently throughout the network. They find that fake news tend to have a lower heterogeneity than true news, due to having less dominant broadcasters. This follows the findings that fake news tend to have less initial sharers and grow over time by depending on branch spreading processes in order to succeed in reaching many individuals. Both of these papers paint the picture of fake news as more virus-like in its behaviour than true news, which tends to stay mostly in the immediate vicinity of the main broadcasters, while fake news weaves its way around the network and its individuals much more deeply.

How can we fight the growing threat of fake news in social media? This is the question at the heart of a field of study with high current interest.

Some use data-mining techniques to detect fake news [37], others use the parameters of rumour spreading models to identify news vs rumours [17]. Here we intend to take advantage of the multidisciplinary nature of Network Science, with a focus on the spread of rumours in structural populations. This focus will help us understand how topology affects the spread of misinformation while in the presence of its truthful counterpart (information that is fact-checked and correct which disproves the rumour).

Hartley and Vu [15] propose insights into intervention policies for fighting fake news through an equilibrium model. They conclude that an intervention policy which shifts a digital citizen's behaviour to one of a "high effort" mindset, can make easier the critical evaluation of fake news. Another intervention pol-

icy is to reduce the utility one gets from engaging with fake news. The former empowers the individual's awareness and detection of fake news, while the latter tries to reduce the dissemination of fake news on social networks and social media (also discussed on Lazer et al. [23]). Our focus here will be on the former intervention type, where we try to raise individual awareness and resistance to fake news. This can be done via training, promotion of fact-checking, etc.

Let us consider that misinformation starts from a single individual. Then, several approaches can be considered to describe how to prepare the population against it. Here we consider two.

Top-down. For this approach we can picture that a government or company selects a fraction of the population and actively influences them with truthful information, increasing their awareness to resist the false information. In a realistic situation, these individuals would then share and spread the information. But, in that situation, not all would blindly believe the information. In this approach we consider individuals that blindly believe the truthful information they receive and that they do not spread it. This is similar to classical vaccination, where a fraction of the population is immunised against a virus. This way, informed individuals will only act as barriers to the spread of misinformation, and will not enter in dialogue with both misinformed individuals and ignorants regarding this matter. This approach can be defined as an external intervention.

Bottom-up. Here we consider coupled dynamics with competitor rumours. We can see this as a sort of viral vaccination, where we can think of truthful information, or the awareness of the fake news, as a virus. Much like misinformation (which behaves as a competing virus), individuals who know or are aware of the truth will share it with others. Sharing occurs through social interactions, which leads to spreading throughout the network. For easier comparison in how the two "viruses" (real vs fake information) spread according to a model, only a single individual starts the truth chain. This approach can be defined as an internal, self-organised, intervention.

Both of these approaches can help us see how three different networks change how both misinformation and awareness spreads. For simplicity and as an analogy to epidemiology, the individual who is aware of the truth will be referred to as vaccinated, and awareness or awareness spreading referred to as vaccination. These networks are Barabási–Albert [3], Dorogovtsev–Mendes–Samukhin [7], and a specific subset of Lancichinetti–Fortunato–Radicchi [22]. Their structures, low clustering, highly clustered hierarchical, and community, respectively, give three different points of view for information spreading and how each network properties affect it. Does having an hierarchical community make it more susceptible to attacks of misinformation, especially targeted attacks? Can we combat it through a stronger and faster vaccination? Do communities prevent misinformation from spreading or do they facilitate it? What is the polarisation of opinions within these structures?

To try to answer these questions we extended a rumour spreading model to study various situations regarding how misinformation and awareness starts and spreads. In order to do this, an effort was made to develop an open-source system to simulate propagation models with a focus on modularity, including models, networks, strategies and data processing techniques.

Another aspect that is studied is vaccination stifling. Vaccinated individuals can stifle[1] neighbouring

[1] Not a commonly used verb, but is used for historical reasons.

misinformed individuals into stopping their spread. Before, whether spreading or not, vaccinated individuals simply acted as uncontactable barriers to misinformed individuals. If we imagine a social network where we can see other's beliefs, maybe through their description or post history, we can consider that without this stifling, vaccinated individuals would simply ignore, or block out, any contact with people that express or have expressed the false opinion. This way misinformed individuals could not contact them. With stifling, anytime a misinformed individual contacted a vaccinated one, the vaccinated would interact back and try to change their mind in order to stop them from spreading the false information.

What was studied here can also be examined with epidemiological models. In Appendix A we share a previous network science project. It studies both BA and DMS networks under an epidemiological model against various intervention strategies. Some of those findings served as a precursor for the present study, using most of the same networks and having a focus on topology.

This dissertation's structure is as follows: On Section 2, we give an overview on networks, spreading models with a focus on rumour models and intervention strategies. On Section 3, we discuss the system to obtain our results, our rumour model used and the networks chosen for study. On Section 4, we discuss our external intervention (top-down approach) results. On Section 5, the internal self-organized intervention (bottom-up approach) results are addressed. On Section 6, we finish with our concluding remarks on the research and possible future work from questions that emerged during the research period.

2 | Key Concepts and Related Work

2.1 Networks and their Properties

Networks, or graphs, are composed of nodes and edges. Nodes can be seen as individuals, and edges are links between these individuals (see Fig. 2.1a). All nodes that are connected from a given node, are called its neighbours. Networks can serve to model a multitude of problems in various areas. While connections between individuals can be made to have only one direction (see Fig. 2.1b), in this thesis we will focus mainly on undirected networks. Before we talk about various types of randomly generated networks, we must first discuss some properties of networks.

This is a quick overview of necessary and interesting network properties to keep in mind when comparing results between different network types as they might be affecting them.

Degree k. Firstly, the degree of a node n is the number of connections that it has with other nodes and the average degree of a network <k> is given by $\sum_{i=1}^{N} k_i/N$.

Degree Distribution $P(k)$. Given the degrees of the nodes on the networks, we can also compute the degree distribution, which is the probability distribution of said degrees in the network. Networks are considered to be scale-free when their degree distribution follows a power-law, $P(k) \approx k^{-\gamma}$, with $2 < \gamma < 3$.

Clustering Coefficient C. This quantifies how many neighbours of a given node know each other. If we imagine the network of a village and a city, where in nodes are the individuals and the connections exist between nodes if they know each other, then in a village the clustering coefficient would be higher than in the city because the tendency is for the villagers to all know each other, while in the city people are more disconnected.

Average Path Length APL. Another measure, the average path length APL, is the average length of the shortest path between all pairs of nodes in the network. The length of a path is the number of connections taken. If a node is the neighbour of another, then the shortest path length of that pair would be 1. If they where the neighbour of a neighbour instead, then it would be 2, and so on.

Another set of measures are the centrality measures, which help us identify some of the most important nodes in the network, for a given interest. We will talk about four of them, which will be relevant later.

Degree Centrality $C_D(n)$. This is the number of incident links to the node n. In our case of undirected networks, it's all the edges of the node.

Closeness Centrality $C_C(n)$. It quantifies how close a node is to every other node in the network. It's calculated by the average path length of the shortest paths between the node n and every other node. The lower it is, the closer a node is to everyone else.

(a) (b)

Figure 2.1: Two examples of networks, where nodes are represented by circles and edges are the lines between them: (a) is a network with undirected edges and (b) has directed edges.

Betweenness Centrality $C_B(n)$. It's a measure of how many times the node n is part of the shortest path between all pairs of all other nodes. If we have two communities where a single node acts as a bridge between them, then that node will have a high betweenness centrality since for a message to get from one community to the other, it must pass through that node, making it an important and central node in that network.

Eigenvector Centrality $C_E(n)$. Lastly, this centrality measure indicates how influential the node n is. It assigns a value to each node n based on their links to other nodes and their values, where links to high valued nodes contribute more to the value of node n, than low valued ones. It captures the tendency for highly influential people to be connected to other similarly influential individuals.

In this thesis, we only used Degree Centrality (equal to Degree due to undirected networks) as a measure to find the most important nodes in a network, but the other four are also possible choices.

2.1.1 Network Types

When regarding epidemiological and rumour spreading models, we want to generate random networks in order to get statistical significance in our results. The following are some models/algorithms to randomly generate networks.

Erdős–Rényi (ER). This random graph model was first introduced by P. Erdős and A. Rényi [31]. Given n nodes and M edges, we choose one of the possible graphs that can be generated using that configuration, as seen in Fig. 2.2a. A variant of this model, which is currently more used, is the one proposed by Edgar Gilbert [12], where the graph is constructed by connecting the n nodes randomly given a probability.

Watts–Strogatz. Due to the ER networks failure to capture the properties of real-world networks, Duncan J. Watts and Steven H. Strogatz [43] proposed a model which included small-world properties such as a low average path length and high clustering. The model first constructs a regular ring lattice, where each node is connected to K neighbours, and then for each edge, rewires it with a given probability.

Barabási–Albert (BA). Another limitation of ER networks was the lack of a power-law degree distribution, which accounts for the emergence of hubs (nodes with very large degrees), leading to a class

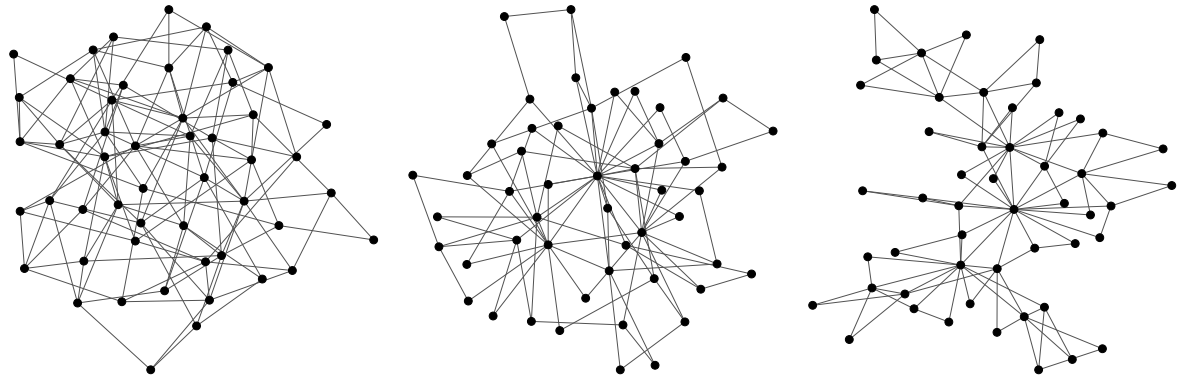

(a) (b) (c)

Figure 2.2: Example of some network types: (a) Erdős–Rényi, (b) Barabási–Albert, and (c) Dorogovtsev–Mendes–Samukin.

of networks called scale-free (SF) networks. Albert-László Barabási and Réka Albert [3] proposed a new model which would deal with this limitation. This new model consists in creating a fully connected network of m_0 nodes and then adding new nodes, one at a time, where the new node is connected to other m nodes. This means that the older a node is, the higher the chance of having a high degree. This is called preferential attachment. An example of a BA network can be seen in Fig. 2.2b.

Dorogovtsev–Mendes–Samukhin (DMS). Another scale-free network model is the one proposed by S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin [7], but different from the latter due to it generating hierarchical networks as seen in Fig. 2.2c. It starts with a fully connected graph of 3 nodes (a triangle), and at each time step, randomly selects an edge and adds a node connected to the two nodes of that edge.

Lancichinetti–Fortunato–Radicchi (LFR). A community network model used to benchmark community finding algorithms proposed by Andrea Lancichinetti, Santo Fortunato and Filippo Radicchi [22]. A myriad of parameters such as the number of nodes, degree distribution exponent, community size distribution exponent, average degree, minimum community size, and much more, allows us to generate community networks tailored to our needs. The algorithm starts with N isolated nodes and assigns each node to a community of a size given the power law distribution using the community size distribution exponent parameter. Each node is also assigned a degree according to the power law distribution using the degree distribution exponent parameter. Each node, per community, is given an internal degree, a subset of its degree, for the intra-community edges. The remaining edges connect the node to nodes of other communities. Following these internal degrees, nodes within the same community are attached, until all internal degrees are fulfilled. After, we do the same with the external degrees to nodes in other communities.

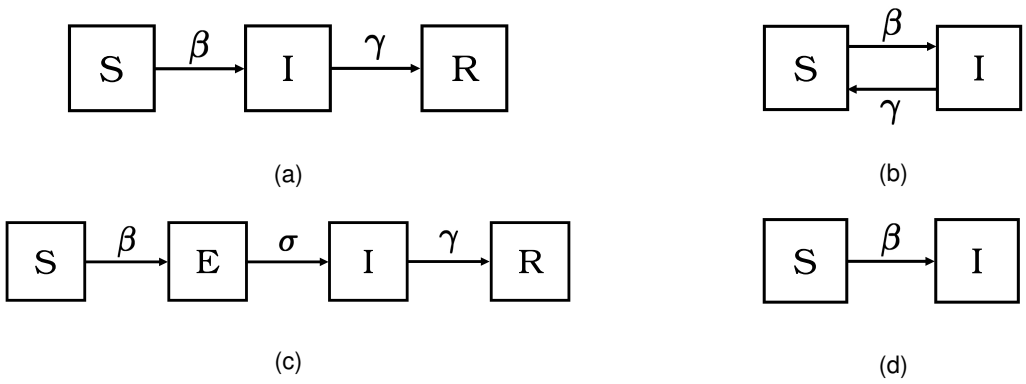

Figure 2.3: The various classical epidemic models with β as the infectivity rate, γ as the recovery rate and σ as the incubation rate. (a) SIR model; (b) SIS model; (c) SEIR model; (d) SI model.

2.2 Single Spreading Models

In this section we will give an overview on some single virus, rumour or idea, spreading models. For easier reading and consistency between models, some of the parameter symbols were changed from their original sources.

Classical Epidemic Models.

From the pioneers of modern epidemiology, Ross, Hudson, Kermach and McKendrick, SIR was born in order to try to model an epidemic spread process in a population (see Fig. 2.3a) [34, 35, 19]. It is a compartmental model which separates the population into susceptible (S), infected (I) and recovered/dead (R), states. The population starts fully susceptible, and a few are selected to be infected. Through contact infection, susceptible individuals are infected and in turn they can infect others through contact. This contact is mediated through an infectivity rate, leading to the success or failure of the infection through contact. Infected individuals can also stop spreading the disease by recovering from it (gaining immunity) or by dying. This transformation is also mediated by a recovery rate, but does not require contact. At any time step an infected individual can recover given this rate.

There are modifications to this model which lead us to: (1) SIS model (Fig. 2.3b), where there is no recovery/death and the infected return to susceptible when they recover (modelling infections which do not have long-lasting immunity); (2) SEIR model (Fig. 2.3c), where there is an intermediate exposed (E) state between susceptible and infected that represents incubation; and (3) SI model (Fig. 2.3d), where there is no recovery or return to susceptible.

Classical SIR Rumour Models.

Following up on a comparison between virus spreading and rumour spreading, Daley and Kendall [6] proposed a rumour model, where there are ignorants, spreaders and stiflers. Ignorants are individuals that have not heard the rumour. Spreaders are "infected" with the rumour and spread it through contact. Stiflers know of the rumour, but no longer spread it. It introduces a decay to the spreading process

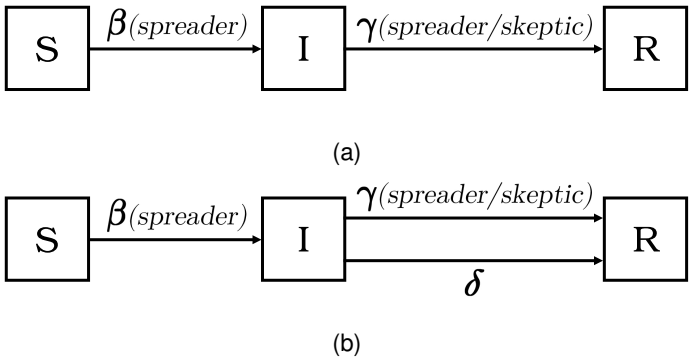

Figure 2.4: (a) The Maki-Thompson model (MK model), with ignorants (S), spreaders (I) and skeptics (R). For the infection rate β, ignorants become spreaders when they come in contact with other spreaders. For the stifling rate γ, spreaders become skeptics when they come in contact with a spreader/skeptic. (b) The Nekovee et al extension of the MK model through the addition of a forgetting rate δ.

where the rumour is forgotten or loses its novelty value. This decay happens through the interactions between spreaders and stiflers or two spreaders. In both cases, the spreaders become stiflers.

The model was further improved upon by Maki and Thompson (MK) [26], by making the rumour spread take place with directed contacts, and in the case of rumour decay, only the spreader that initiated the contact becomes a stifler. The dynamics of this MK model on complex heterogeneous networks were studied in Moreno at al. [28].

In Nekovee et al. [29], the MK model was improved to include a forgetting mechanism. This mechanism implements the realistic possibility of a spreader forgetting the rumour without interaction. This and the stifling mechanism from MK happen at their specific rates. Nekovee et al. found that this model showed new behaviour that was not present in simpler models. A critical threshold for the spreading rate, where the rumour cannot spread in ER networks, was found. A similar but much smaller threshold was also found for finite-size scale-free networks. Moreover, in SF networks the initial spreading rate was much higher than in ER, revealing their propensity to spread rumours, in a way similar to viruses.

SIZ, SEI and SEIZ.

In trying to model the spread of an idea, Bettencourt et al. [4] applied several novel models to the spread of Feynman diagrams in various scientific cultures. SIR was considered as a baseline model, as it has decay of adopters through recovery. But Bettencourt et al. found there is a fundamental flaw in it, where recovery or immunity have no meaning in idea spreading. Thus, such a state is removed from the next models and, in turn, they introduce population losses at a given rate.

The effects of idea spreading that were sought out in Bettencourt et al. were adoption through multiple contacts, competition and incubation. Competition appears by reasoning that individuals can oppose an idea once they've come in contact with it. With this mechanism, the model SIZ was made, where Z is the skeptic state. Susceptible individuals can now become either infected or skeptic, depending on

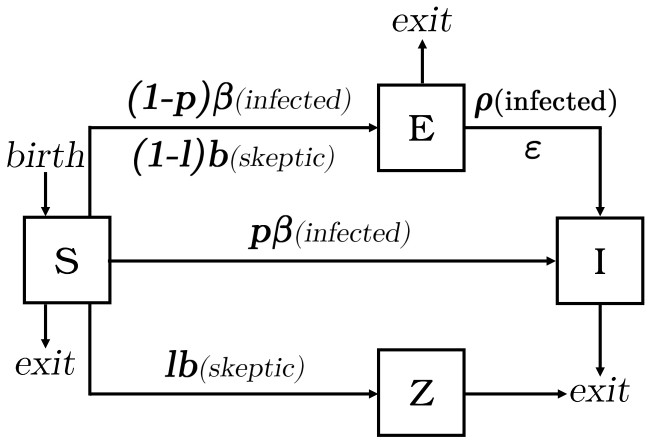

Figure 2.5: The SEIZ model. The parameters shown are as follows: β is the S-I contact rate, ρ is the E-I contact rate, b is the S-Z contact rate, ϵ is the incubation rate, l is the sucess probability of a Z turning an S into Z and p the probability of an S being infected imediately. Population dynamics are shown as birth and exit, and are regulated with their own rates.

who they contact with. The problem with this model and the SIR model is the immediate transition to infected through contact. This transition is not realistic and there was a need to introduce some latency. In order to introduce this latency, an incubation state E was added to the idea spread in SIZ. When susceptibles come in contact with an infected, they can immediately accept the idea or enter the incubation state, where they might later adopt the idea. This also introduces a side effect to the contact between skeptics and susceptibles, where susceptibles become incubated instead of skeptics. This new state and transitions lead to the SEIZ model.

Another model SEI, which is a subset of SEIZ with only incubation, was also tested. While fitting the various models to the Feynman data, SEIZ was found to have the lowest of the smallest average deviation, but most of the improvement came from SEI, when compared to the models without latency such as SIR. This improvement came from a better fit of the early stages in the spreading process.

It should be noted that SEIZ assumes there is an initial population of skeptics, since the only way to become a skeptic is to come in contact with one. The "good idea" that is stated in [4] can be seen as the reason for this, because a "good idea" will tend to surely infect (E or I) and not cause someone to reject it altogether (Z) without there being someone who already rejects it. In the case of the study, these were those who wanted to maintain the status-quo regarding the diagrams used.

In Jin F. et al. [17], SEIZ was used to model the spread of true news and rumour stories in Twitter and was accurate in modelling the information spread. A method for rumour detection using the parameters of the model was also shown to be an accurate tool for identifying true news vs rumours in some cases.

SIHR.

In Zhao et al. (2012) [47], the MK model is extended by introducing forgetting and remembering mechanisms into an hibernator state H. An ignorant individual can also become a stifler, when exposed to the rumour. These extensions lead to the SIHR model. Spreaders can forget the rumour and hi-

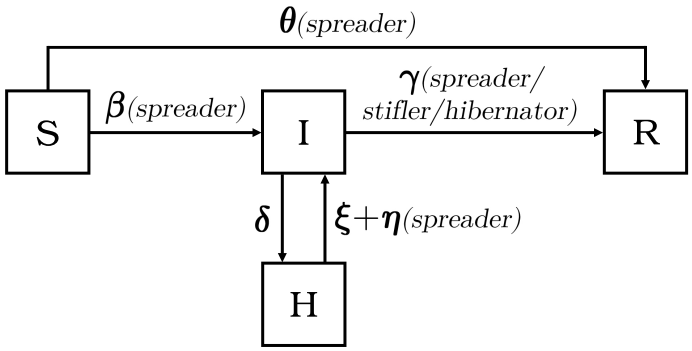

Figure 2.6: The SIHR model. The parameters shown are as follows: β is the infectivity rate, γ is the stifling rate, θ is the refusion rate, δ is the forgetting rate, ξ is the spontaneous remembering rate and η is the remembering rate.

bernate. Hibernated individuals can become spreaders, either spontaneously or through contact with a spreader, at their respective rates. These hibernators reflect the repeatability of rumour propagation, what Zhao et al sought to test.

Through numerical simulations, Zhao et al. (2012) found that the direct transition from ignorants to stiflers accelerated the rumour's end, while lowering it's peak influence. It was also found that the forgetting and remembering mechanisms reduced the peak influence, but postponed the rumour's end. Furthermore, the existence of hibernators was entirely ruled by the remembering rate.

SIR with variable forgetting rate.

In Zhao et al. (2013) [48], the MK model is extended to address the use of a variable forgetting rate, where the longer an individual knows a rumour, the more they are bound to forget it.

The spreader state is divided into spreader and spreader* states. The first works as in the classical SIR model, where ignorants that are infected go into this state, and when spreaders of this state contact another spreader/stifler they can become skeptics at a given rate. If a spreader fails to contact another spreader/stifler or they fail to become a skeptic given this contact, then they become a spreader*. The latter can forget the rumour and become skeptic, at a variable rate that depends on the spreader's age since first being infected. If he fails to become a skeptic, the spreader* becomes a spreader, again. The variable forgetting rate is modelled after an exponential function.

The model was applied to the platform LiveJournal where it was found that increasing the initial forgetting rate, or the faster the forgetting rate increases over time, causes a smaller final rumour size. Moreover, the final rumour size is larger when a variable forgetting rate is used, rather than a constant one.

SCIR.

Xiong et al. [44] proposed an information diffusion model based on the information sharing process in websites like Twitter, where there are susceptible, contacted, infector and refractory states. Susceptibles

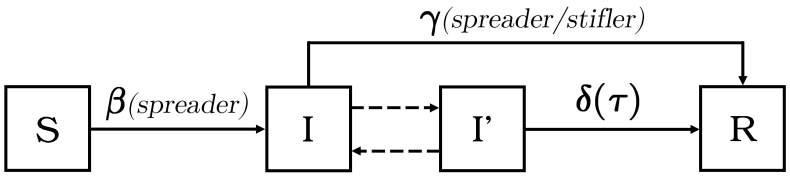

Figure 2.7: The SIR with variable forgetting rate model. The parameters shown are as follows: β as the infectivity rate, γ as the recovery rate and $\delta(\tau)$ as the variable forgetting rate where τ is the time steps passed since the individual became a spreader. The dotted lines between spreader and spreader' states represent the intermitent switching between them (see model's section).

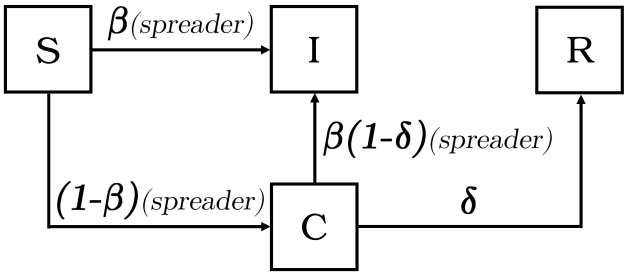

Figure 2.8: The SCIR model. The parameters shown are as follows: β is the infectivity rate and δ is the spontaneous refactory rate.

are ignorant of the topic. Infectors spread the topic to their neighbours. When users have learned of the topic, but choose not to spread it, they are in the contacted state. Contacted users can become disinterested and change into refractory users, which know of the topic but will never share it again. These users can also become infectors when contacting an infected, at a given infection rate. Infectors and refractory agents are absorbing states.

On SF networks, Xiong et al. found that the higher the connectivity, the higher the proportion of infected agents and the smaller the number of refractory agents, given that contacted agents are infected more often. This proportion of infected agents levels off, above an average degree of 30. Given this, there are always refractory agents who "refuse" to spread the information.

SEIR(Education).

Based on the studies in [26, 29], Afassinou et al. [1] developed a rumour spreading model around the idea of the education rate of the population. It is assumed that a more educated individual has less propensity to believe rumours, due to their ability to assess rumour credibility, when compared to someone less educated. This education rate will then influence the forgetting and stifling mechanisms.

The model separates the ignorant state into educated (E) and uneducated (I) ignorants, while keeping the spreader and the stifler states from [29]. This modification gives us the SEIR model. When ignorant individuals contact with spreaders, they can choose to spread the rumour and become spreaders, or not to spread it and become stiflers. Educated ignorants will have a harder time accepting a rumour

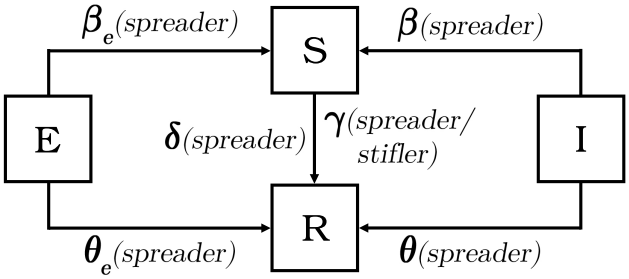

Figure 2.9: The SEIR(Education) model. The parameters shown are as follows: β is the infectivity rate for the uneducated (β_e for educated), θ is the rejection rate for the uneducated (θ_e for educated), γ is the stifling rate and δ is the forgetting rate. Educated individuals will believe rumors less and we will have $\beta_e < \beta$ and $\theta_e > \theta$.

without questioning its validity, becoming spreaders or stiflers at lower and higher rates, respectively, when compared to uneducated. The process of a spreader losing interest and becoming a stifler is maintained, as is the process of contact spreader$-$spreader or spreader$-$stifler causing the initiating spreader to become a stifler. In this way, the education rate should have an impact on how many ignorants become spreaders.

Afassinou et al. found that a rumour's final size decreased as the education rate increased, revealing its role in spreading, on top of the stifling and forgetting mechanisms introduced in [29]. In their numerical simulations, education rate was found to lower the maximum peak of spreaders, while also delaying its occurrence. As this peak is reached and the density starts to dwindle, the density of stiflers quickly rises and the rumour dies out. Furthermore, the forgetting rate was found to impact the rumour's final size by just 4%, but this mechanism should not be neglected in the case of rumour spreading.

Fact Checking Model.

In Tambuscio et al.(2015) [39], the role of fact-checking and its ability to completely remove a hoax from a network, was studied. In their model, there are susceptible, believers and fact-checkers. Susceptibles can become believers or fact-checkers depending on the rate of spreading, the credibility of the rumour and the state of their neighbours (believers or fact-checkers, respectively). Believers and fact-checkers can also forget the hoax and become susceptible. The fact-checking happens spontaneously given a probability to verify, from believer to susceptible.

What Tambuscio et al. found was that the fact-checking probability can always remove a hoax no matter the credibility. The threshold value for this complete removal depends on the credibility and the forgetting probability, but not on the spreading rate. These results were obtained for BA, ER and real networks. Though the model tends to behave similarly in BA and ER networks, under certain conditions the hoax is removed from the ER network, or becomes endemic (only fact-checkers) in the BA network.

Building up from [39], Tambuscio et al. (2018) [38] tested the role of network segregation in the spreading of a rumour. This segregation divided the network into gullible and skeptic groups, and was associated with a segregation rate *s*. The higher *s*, the more intra-group and less inter-group connections

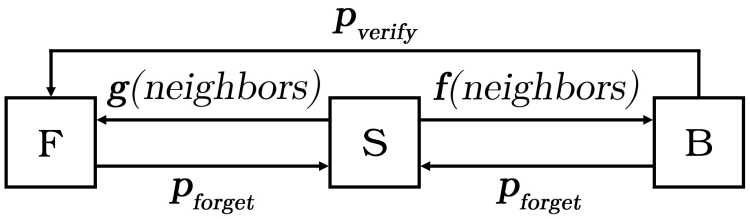

Figure 2.10: The Fact Checking Model. The compartments are the susceptible (S), the believers (B) and the fact-checkers (F). The parameters shown are as follows: p_{verify} is the verifying rate, p_{forget} is the forgetting rate and f and g are spreading functions which describe the hoax and its debunking spreading process, respectively. Both functions take into account the neighbors of the individual, a spreading rate and the credibility of the hoax.

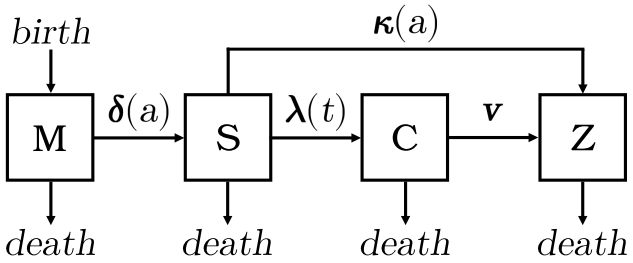

Figure 2.11: The Age-structured model. The parameters shown are as follows: $\delta(a)$ is the age-specific rate to leave the M state into S, $\lambda(t)$ is the force of infection, $\kappa(a)$ is the skepticism rate and v is the recovery rate. Population dynamics are shown as birth and death,and are regulated with their own rates/parameters.

the nodes had. This means that with higher *s*, the skeptics are more likely to connect with other skeptics than with gullible nodes. The same applies for gullible nodes.

When *s*, or the gullibility of the gullible group increase, the rumour's final size tends to increase. However, the forgetting probability plays an interesting role. When it's small, increasing the segregation increases the rumour's final size. But segregation has no effect when the forgetting probability is greater, with the rumour's final size tending to be larger in this case.

Segregation can influence the rumour spreading differently, depending on the forgetting probability. When forgetting is less probable, gullible individuals are more influenced when there is less segregation by being in greater contact with skeptics. And when forgetting is more probable, less segregation leads to skeptics being more influenced by gullible believers.

Age-structured epidemic models.

Noymer et al. [30] present two epidemic models on an inquiry about the "urban legends" type of rumour spreading, that seem to perpetually plague younger age demographics.

Individuals are born into an M state where they cannot be infected as they are too young to under-

stand and learn of these rumours. They leave this state at an age-specific rate, which is modelled with a delayed exponential decay, into the susceptible group, S. From here they can be infected or become skeptics. Infection is ruled by a force of infection, which takes into account the proportion of infected agents, state C, and a constant that represents the population mixing and probability of transmission on contact. They can also become skeptics, state Z, with another delayed exponential decay, at a scepticism rate. This rate is delayed based on a given age after which the agents are far less gullible to believe these rumours. The infected agents, for a given recovery rate, can also become skeptics, modelling the loss of interest. Mortality occurs in all states at the oldest age, another parameter.

Initially, the model is run without rumour transmission in order to reach a population equilibrium. In this initial equilibrium, the population follows an acceptable distribution by age, where most susceptible individuals are of a younger age. Then, rumour transmission is turned on and tested for two different recovery rates (0.2 and 0.04). In both cases there are recurring epidemics, with the first being the largest ($\approx$3% and $\approx$4% of infected, respectively) and following an exponentially reduction in size for the next one. For the higher recovery rate, between epidemics, the number of infected drops to negligible levels. On the low recovery rate, it becomes significant, following the second epidemic. In the lower recovery rate, the periods between epidemics show a trend for a non-zero endemic population. With regards to mean age of infection, it tends to start out at a higher age, 16+ years, and then settles at 13 over decades. The standard deviation also settles over time.

In another model, based on the previous, the recovery rate is changed to a force of scepticism. This is to model the believer being challenged by a skeptic and not simply recovering from their belief on their own. This force of scepticism works similarly to the force of infection, as it relies on the proportion of skeptics, and a different population mixing constant. A parameter q for this contrast is the probability of change from believer to skeptic when contact is made with a skeptic.

The population equilibrium is reached in a first run, same as before. Two cases for q are tested: a value of 0.3 and a value of 0.01. With the first case, the model behaves similarly to the previous model with the higher recovery rate. For the second case, we see an infection of 30% of the population, with a rapid and cycleless convergence to an equilibrium of $\approx$5%. We also see a difference for mean age of infection, which reaches a much lower and stable number, meaning only a narrow range of young people believe the rumour and quickly become skeptic as they age and become replaced by new believers(an endemic state).

SIRaRu.

For studying the immunisation of a population against a rumour, Huang et al. (2011) [16] developed a new model with a state that resembles immunised individuals. This model, named SIRaRu, was extended from the MK model and it separates the stiflers in two, where a stiflerA accepts the rumour but doesn't spread it and a stiflerU doesn't accept the rumour and doesn't spread it. When an ignorant is contacted by a spreader, it can (1) accept the rumour and spread it (spreader), (2) accept the rumour but not spread it (stiflerA) or (3) reject the rumour and not spread it (stiflerU). The transition from spreader to stiflerU remains the same as the MK model, following contact with spreaders or stiflers, but lacking the

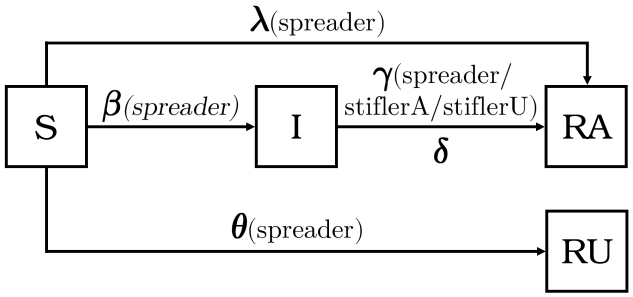

Figure 2.12: The SIRaRu model. The parameters shown are as follows: β is the infectivity rate, γ is the stifling rate, θ is the rejection rate, λ is the acceptance rate and δ is the forgetting rate. We also have that $\beta + \theta + \lambda = 1$.

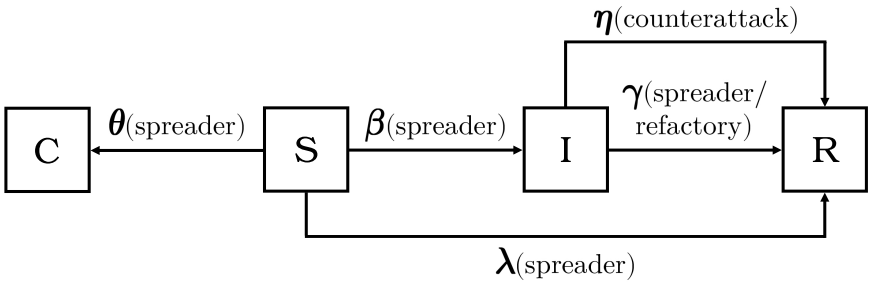

Figure 2.13: The SICR model. The parameters shown are as follows: β is the infecivity rate, γ is the stifling rate, θ is the refuting rate, η is the persuading rate and λ is the ignoring rate.

forgetting/loss of interest. The rate at which an ignorant becomes a stiflerU is based on the credibility of the rumour. The smaller the credibility, the smaller the number of individuals who accept the rumour. The stiflerU is the state that resembles immunised individuals, as individuals who know the truth about the rumour will not accept or spread it, mimicking the expected behaviour of real vaccines where vaccinated individuals will not contract or spread the disease.

The model above was initially used to study immunisation in small-world networks, and was further extended and studied by J. Wang et al. (2014) [42] for homogeneous and inhomogeneous networks. The model was extended to have forgetting/loss of interest as it is more realistic. The effect of the topology was studied and found to influence the rumour spreading process. In BA networks, the rumour spreads faster but reaches a lower maximum rumour influence than in homogeneous networks. It also reaches a smaller final density of stiflerA in comparison, while the rumour termination time remains the same. It was also found that the smaller the credibility, the smaller the maximum rumour influence, and so it becomes a weak point to attack when dealing with immunisation. Following the previous studies, the forgetting and stifling rate were found to reduce the rumour's influence as they increase.

SICR.

With the goal of studying the self-resistance of a network to rumour spreading, Y. Zan et al (2014) [45] proposed an extension to SIR with the addition of an absorbing counterattack state, leading to the SICR model. This state follows the line of thinking that upon hearing the rumour, some individuals might reject it and work towards stopping its spread. This rejection is mediated by a refuting rate θ. A counterattack individual will try to reform its infected neighbours, where they will become refractory at rate η, the persuading rate. It might seem to work the same as the state of stiflerU in SIRaRu but stifling rate γ and η are different rates in this model. Y. Zan et al. mention similar models that do this.

Another model proposed by Y. Zan et al., for comparison purposes, is the adjusted-SICR model, where the stifling process that happens on contact, between a spreader and a spreader/stifler, is removed. This is due to a psychological phenomenon where spreaders may get the idea that there are still individuals interested in the rumour and so they will keep on propagating it.

The model SICR is compared with SIR and adjusted-SICR, and it is found that, in both cases, with SICR we have (1) more susceptible at the end of the rumour epidemic, meaning less people see the rumour and (2) a smaller maximum rumour influence (peak density of spreaders), meaning the counterattack group is slowing the speed of rumour spreading and the density of spreaders. Through more testing of SCIR, they found that increasing θ increases the final density of susceptibles and decreases the peak of spreaders. The same follows for increasing the self resistance of the model (a combination of θ, η and γ). But when comparing the effects of stifling and persuading on the final density of susceptibles, we see a significant increase when increasing stifling versus an insignificant change with persuading. In fact, for fixed values of stifling, increasing the persuading rate leads to a fluctuation of the final density. When comparing them, in regards to the maximum rumour influence, both rates contribute to a significant decrease but the stifling rate shows a more dramatic slope.

Other models.

Other models which are too complex to explain here but are nonetheless notable mentions are (1) the Energy Model by S. Han et al. [14], which takes inspiration from physical theory and focuses more on the microcosmic aspects of rumour spreading; (2) ICSAR (8 state model) by N. Zhang et al. [46], which massively extends the SIR rumour model in order to study rumour spreading mechanisms that are not captured in SIR such as information attraction, subjective and objective identification of rumours, trust, effect of experts and personal preferences.

2.3 Intervention Strategies

Here, we will give an overview of various intervention strategies. We define intervention strategy as a method of applying vaccines to a population and their effectiveness into how well they can contain the spread of the infection. Before going into them we need to first contextualise intervention in the matter of fake news/rumours.

There are many ways to describe what the term fake news means and scholars don't agree on a single definition. Some say fake news are fabricated news that resemble news media content in form, but lack the norms to ensure its authenticity and validity [23]. A similar definition is that they are news articles that are intentionally and verifiably false, and could mislead readers [2]. Edson C. Tandoc et al. [40] mention that the term fake news currently describes false stories that spread in social media and is even used to discredit news organisations. Furthermore, they mention that the old definition was regarding related but distinct types of content, such as news parodies, political satires, and news propaganda.

With these definitions, what can we say about vaccines in this context? If we assume that intervention in a population is a conscious effort from those doing said intervention, and the ones receiving it, then we are looking at informing the population. Giving information to individuals in order to combat the spread. To damper the spread, the individuals we intervene upon need to become adversaries to this fake news. This can happen directly as stiflers or counterattack states, or indirectly as educated individuals. In the first case, we are informing against a specific threat, and in the second we are educating on the subject.

There is also the matter of who provides the intervention. We could have a central entity that provides the service of informing the population of the truth, perhaps with facts to support it. Or, it could also be spread through peers in the population, for instance through word of mouth. In rumour spreading, unlike in virus spreading, and in most of the models discussed here, peer transmitted intervention happens due to the stifling mechanism. While a stifler cannot convert a susceptible into another stifler, it can coerce a spreader into becoming a stifler, essentially spreading the intervention. If we assume a central entity will inform a population and create a "vaccinated" fraction in it, then peer intervention will follow.

Without considering stifling and forgetting mechanisms for a moment, let us consider a network of spreaders, which was initially just one, and spread over time. Each one of those spreaders is linked to, at least, another spreader, from where they got infected. This means, we can start in one of the spreaders and reach any other one, given their shared ancestor. If we now include stifling, and place one stifler connected with one of these spreaders, then they will be able to spread their "vaccine" throughout all the spreaders and make them stiflers, given time. With forgetting added, stiflers can appear in the spreader network and further this intervention. Even if this could be regarded as a self defence mechanism of the network/model itself, we want to minimise how much it spreads. So on top of this inherent intervention, we can add our own, in the form of a central entity driven intervention. By creating more stiflers, or individuals more prone to becoming them like the educated, we can create a sort of immune system inside the network which would fight the "infection" faster and more aggressively. This is different from virus spreading, where vaccinated individuals only act as a static barrier to the infection.

The following are some intervention strategies first conceptualised in epidemiology, yet they should work given their underlying nature.

Random. The most simple intervention strategy to conceive is random intervention. It means to simply intervene randomly on the population until we reach the wanted quantity. It tends to be the worst, because it uses no knowledge of the current spreaders or network topology, or the underlying network properties. Being random, it is a good baseline to compare to other strategies.

Ring. It's another easily conceived strategy, where we intervene on the neighbours of the spreader population in a ring manner in order to contain the "infection". This makes it so it cannot spread in any of the available connections, effectively curbing the spread. In real life social networks, given their dynamic nature, with connections constantly appearing and disappearing, this strategy is applied to the people we know will have the most contact with the identified infected people. An example of this is medical staff. The issue with this strategy is the need to be applied when the infection spread is at its infancy, where it is most effective [13, 20]. When the infection is more spread out, applying vaccines quickly around the neighbours of the infected is more difficult in a real life scenario, making it less effective and more prone to containment failure. But even in some cases early containment may not be possible because we may not have perfect information regarding who is infected [21]. Another problem is the need for topology knowledge, where at infancy it might be easier to acquire, when the spread is larger it becomes nearly impossible. Given full topology knowledge and full knowledge of current spreaders, we could vaccinate all neighbours of all the spreaders and the spreading would stop.

Hub. A method which exploits the topology properties and how viruses spread in scale-free networks (also called targeted immunisation) [32]. In these networks, hubs are the most important aspect of transmission of viruses, rumours and even messages, due to being nodes with high degree and centrality. If they become infected, they can spread the infection to a great number of nodes. From this, it seeks to vaccinate these propagation highways in order to diminish the spread and its speed. The usual hub method is to intervene given the degree, starting at the highest. The problem with intervening on hubs is the need for perfect topology knowledge in order to ascertain which nodes are hubs.

Hubs can be seen as more than just high degree nodes. As discussed in [9], high eigenvector centrality might prove more appropriate, because being a hub would mean to be connected with other hubs, making it an important node in the network with regards to propagation (of information, a virus, etc). Even betweeness or closeness centrality can be used, as they indicate that the node is important to propagation inside the network.

Acquaintance. Contrary to the latter, this is a method that requires no knowledge of the network's topology [25, 24, 5]. This method selects, randomly, a fraction of the population and for each one selects randomly an acquaintance (neighbour) and intervenes on it. Doing this, there is a higher chance of selecting hubs because randomly chosen nodes are more likely to be an acquaintance of a hub given their high degrees, and so it nears the Hub method. Not requiring knowledge of the topology makes it a fairly cheap strategy to implement.

Gallos et al. [11] discusses various improvements to this strategy. The first is to ask the individuals from the first random selection to point us to one acquaintance who is more connected than them and vaccinate them. If there is no such friend, we continue to the next random selection. This method proves to be more effective than the basic Acquaintance method. Also, a second variation is studied, called the enhanced acquaintance immunisation (EAI), which proves to be even more effective, and follows closely the results of Hub immunisation. In this variation, we ask the individual to point us to a random friend who has degree larger than a certain threshold k_{cut}. If there is one, we vaccinate him, otherwise we continue with the random selection. These changes increase the chances of selecting a hub.

3 Proposed Approach

In the following sections we will discuss the software solution for generating simulation data, our rumour spreading model and the networks selected for study.

3.1 Software Solution

When considering how to run the simulations, we figured it would be helpful to build a modular system to facilitate the results generation for contact/propagation models. The system would allow the switching of models, networks, strategies, parameter values, data processing, etc., easily. Also, these modular parts were simple to start working with while having enough depth so experienced users could tackle more unique problems.

The implementation of this system was done in Python, using the library *networkx* for graph handling. The software is open sourced and available at this repository.

The system is split in 6 main parts: Networks, Models, Strategies, Parameters, Simulator Core and Data Processors.

Networks

Networks are simply defined by their graph generation. There is a base Network class that serves as an interface between our code and the *networkx* graph, by implementing auxiliary functions to help the simulator. When we want to add a new network, we extend this class and define the graph generation function. This function expects a *networkx* graph, whereas the logic for adding nodes and edges is left to the algorithm.

Models

Models are defined by 4 things: States, Parameters, Transitions and the ending condition.

States are a list with the names of each state in the model.

Parameters are a dictionary which associates a name with a value (e.g. β and 0.2).

Transitions are a dictionary of transition names (e.g. 'IS', explained further down) and their functions. There are two kinds of transitions: (1) pair and (2) self. (1) Pair transitions are transitions between two nodes, with certain states, that requires interaction between the two. (2) Self transitions are when a node can change its own state by itself with no outside interaction.

When we create this dictionary, the transition is defined as a string of the initiator and receiving state, in that order, where the initiator state is the one that initiates the contact and the receiving state receives it. For example, in SIR, we would have 'IS' regarding the infection transition by the infected I. For self transitions, this string will only contain a single state. These transitions are then associated

with a function that will enact the transition logic. These functions receive the nodes, their states, and the network as input, and output the list of new states (or just the new state for self transitions). The functions can also retrieve the values of the parameters and use them, along with the input, to decide on the outcome of the transition. With the network input, we can also implement complex contagions, since we can check the states of the neighbours of the nodes involved. The output can be the same states as received in the input if no change is needed (e.g. when a transition fails), and no network update will occur. Given the implementation in Python and the use of dictionaries which cannot have repeated keys (transitions), if there are more than one transition between two states, then the logic to handle both will have to be put in the same function.

And, the **ending condition** is when to stop the simulation. It is done via a function that is called at the start of every time step and it receives the step and the network. With these two arguments we can stop the simulation for various possible cases: when a certain time step is reached, when the amount of nodes for a certain state reaches a specific value or when the variation in the network reaches a stable point, etc.

Models also have extra functionality to be used, if needed. What states are considered for transitions for each time step, how to initialise the network in terms of percentages for each state, and what states should be considered for each type of strategy. The first is especially important since it lets us do things like timescales where we can force that only transitions for state A are considered at time step 1, and then only for state B at time step 2, etc.. This will be important when we discuss the models used in this work. The transitions that are considered are only the ones where the state is the main actor, so pair transitions where it's the initiator and all of its self transitions. The user can add their own logic as to what happens for each time step. It is also possible to do more complex patterns since the function regarding this functionality returns a list of states.

Strategies

Strategies are changes to the states of nodes that are applied after the generation of the networks. Strategies can be supplied to the simulator core with an associated purpose, which would then be used to get what states are the nodes changing to (defined in the model). These strategies receive the network and parameters on how to change it (number of nodes to change) and apply their changes directly to the nodes in the network.

Parameters

Parameters are used to supply values to the models. When we are running simulations, it is usual to want to cover ranges of values for some parameter in order to get data regarding certain processes. To cover these ranges, we need to do combinations of all possible values that are requested. If we fix the value of β to 0.5, but want to do θ and γ from 0 to 1 with a step of 0.1, then we will have 121 combinations. In order to simplify the process of passing values or lists of values, there is a whole system regarding the return of all the combinations given the parameter settings passed. From the user point of view, create objects of certain classes like IntervalParameter if they wish for all the values in an

interval with a given step. These will then be passed to the parameter system which will go through and feed all the combinations to the simulator.

Config

The config is a list of settings that include the Parameters and are passed to the simulator. Besides the Parameters, these settings will be fed to other parts of the system where each part will retrieve what information they need from it. However, they tend to be mostly used for the Data Processors which will be discussed further down.

Simulator Core

The simulator core is where the simulations are run. Given the parameter definitions that we received via the Config, the networks and the different sizes of networks, we do simulations for each combination of all of these. The number of simulation repeats for each combination is defined by 'steps' in the Config.

Before starting a simulation, we create the network and initialise the states according to the model definition. Then, given the model, we create a list of all the possible transitions for the current state of the network, both self and pair transitions, kept indiscriminately together in the list. Another list that is kept contains all state changes for the current time step. Then, for each time step, we randomise the transition list, and go through it in order. If the transition is successful, then the new state changes are saved in the changes list. However, before doing a transition we need to check if the node(s) in it have not changed state already, through the use of the changes list. If they have been changed, we skip the transition. After going through all of the possible transitions, we update the network according to the state changes. We also update the transitions list given the changes. We repeat this for the next time steps, until the end condition of the model is met.

The whole process is repeated given the number of repeats, and then again for the next combination of parameter values, until we run out of combinations.

Data Processors

The data processors are modules have access to the simulator's state at various defined steps. They are defined according to a fixed interface. These steps are defined by hook points, which are points in the simulation where specific functions in the data processors are called. For example, when we select a combination of values before and after doing the simulations. This lets an outside system see the actual state of the simulator when it calls the hook points. Before starting the simulation, we can feed it one or more data processors (which are processed in LIFO order). These data processors can range from saving the state (or measures) of the network at each time step, or only at the end; to visualise the simulation in real time; or other uses the user might need. To implement a new data processor, the user has to extend the Generic class and implement the functions for the hook points that they need. The current hook points are as follows:

- Before & After every time step of a simulation

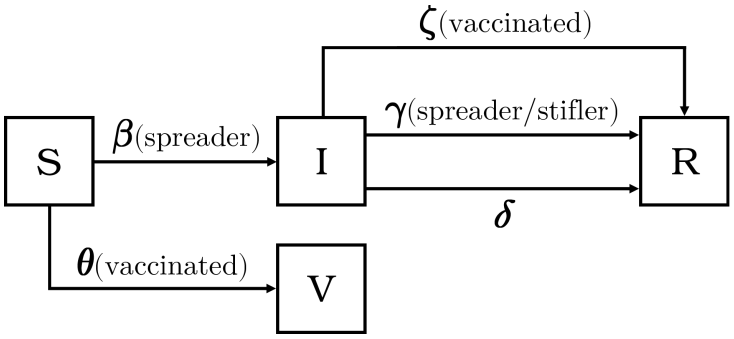

Figure 3.1: Diagram of the SIRV extension model of the Nekovee et al model [29]. The parameters are as follows: infection rate β, stifling rate γ, forgetting rate δ, vaccination rate θ, and vaccination stifling rate ζ. The four states represent: ignorants (S), spreaders (I), stiflers (R) and vaccinated (V).

- Before & After every repeat of a simulation
- Before & After every set of repeats
- Before & After every combination of values for the parameters

There is more that can be added. As the system evolves so will the number of hook points.

3.2 Rumour Model

In order to study awareness spreading (referred to as vaccination from now on) with rumour models we decided to focus on a simpler SIR model, the Nekovee et al. extension of the Maki-Thompson model [29] with two further additions to it. The first accounts for vaccination, in order to facilitate the study of certain behaviours (less degrees of freedom). This addition comes as a new state for vaccinated (or informed) individuals. The second accounts for an interaction between vaccinated and spreader individuals. In the resulting model, we get four states: ignorants (S), spreaders (I), stiflers (R) and vaccinated (V).

Both I and V can convince S to turn over to their side, but not each other. Once an individual believes in a certain idea he will never change to another. While this does not match reality, the simplicity of this model helps in establishing the relationship between the remaining variables.

The vaccinated-spreader behaviour results in a new transition where vaccinated individuals can contact misinformed neighbours and turn them into stiflers. To this stifling behaviour, we'll refer to it as active vaccination. In this addition, vaccinated individuals actively try to convince others around them to stop spreading misinformation. While in reality this could also lead to the others becoming vaccinated, we simplified by not considering this. If they are "convinced", they become stiflers. Furthermore, without this stifling, vaccinated individuals are passive in this regard as they will ignore all contact with spreaders.

Both of these extensions can be made in the same model, which we'll call SIRV, as seen in Figure 3.1. The transitions between states are as follows:

β: S$\rightarrow$I. When an I individual interacts with an S, the ignorant can become "infected" with probability β.

θ: S$\rightarrow$V. When a V individual interacts with an S, the ignorant can become vaccinated with probability θ. The vaccinated state works like a virus. The idea is that this vaccination is the truth opposite to the misinformation spread by I individuals.

γ: I$\rightarrow$R. When an I individual interacts with another I or an R, the spreader can be stifled with probability γ.

δ: I$\rightarrow$R. An I individual can forget the misinformation with probability δ, becoming an R.

ζ: I$\rightarrow$R. When an I individual interacts with a V, it can become an R with probability ζ.

For $\zeta = 0$ we can study vaccination spreading without vaccination stifling. With $\zeta = \theta = 0$, we can study classical vaccination. On Table 3.1 we present the various model parameters for reference.

In Chapter 4, we look at the top-down approach where vaccination does not spread. In this scenario, we could consider states R and V as the same state when vaccination stifling ζ is equal to stifling γ.

Parameter	**Symbol**	**Range**
Number of nodes	N	$\mathbb{N}$
Infection rate	β	[0,1]
Stifling rate	γ	[0,1]
Forgetting rate	δ	[0,1]
Vaccination rate	θ	[0,1]
Vaccination stifling rate	ζ	[0,1]

Table 3.1: Model parameters and their value ranges.

3.2.1 Implementation

When performing the simulations using this model and our software, the simulations are stopped when the number of spreaders (I) becomes 0. While this will not have an influence when we are measuring the amount of misinformation, it can influence the amount of vaccinated individuals which could've kept spreading after misinformation dies out. This will then influence our measure of polarisation (discussed in Chapter 5), since it relies on that amount.

A pseudocode implementation of the SIRV model can be seen in Appendix B.

3.3 Selected Networks

In order to study the behaviours of misinformation spread, and vaccination, we picked 3 scale-free models, BA, DMS and LFR. BA and DMS were selected because they represent extremes in clustering coefficient. BA are low clustered and DMS are highly clustered. LFR on the other hand, gives us a community structure with low clustering. The combination of these three scale-free networks allows us to assess the roleplay of structure, clustering and communities. While LFR is not inherently scale free, it can be configured to generate scale-free networks. For these three networks, we picked a size of 1,000 nodes, which is enough to provide meaningful results, while still yielding fast simulations.

Metrics	BA	DMS	LFR
Degree Distribution Exponent	3.05	2.92	2.69
Average Degree	3.99	3.99	4.33
Clustering Coefficient	low (0.03)	high (0.74)	low (0.14)
Average Path Length	4.08	4.91	6.80

Table 3.2: Metrics for the various networks used. BA and DMS used 1,000 randomly generated networks, while LFR used the 100 generated networks that are used throughout this paper. Degree distribution does not take into account the low and high cutoff points, which would bring the results closer to the expected value of 3.

3.3.1 Barabási–Albert (BA)

BA are scale-free networks with very low clustering (see Table 3.2). Since there are little to no triangles in the network, nodes will not be connected to the neighbours of their neighbours. This will lead to nodes being connected to other nodes outside of their immediate vicinity, which results in a lower Average Path length (APL). This is expected to facilitate both the spread of the misinformation and vaccination. For an example of a BA network with 1,000 nodes, see Figure 3.2a.

3.3.2 Dorogovtsev–Mendes–Samukhin (DMS)

DMS on other hand provides us with a highly clustered scale-free network. Its network algorithm starts with 3 fully connected nodes (a triangle), then randomly selects an edge and adds a node connected to the nodes of this edge (adds another triangle). It repeats these last 2 steps until we reach the number of desired nodes. This creates a highly clustered network which has an hierarchical structure, as can be seen in Figure 3.2b. This topology has interesting properties for information propagation since it resembles a tree like structure. Given the starting triangle (the core nodes), every node that branches off an edge of this triangle will never share a direct edge or non-direct path (not passing through the core) between them. This forces a funnel of information to pass through at least one of the core nodes. If we block any core node, information can still flow freely because of the core triangle. If we block two core nodes, then all of the branches for that core edge will become disconnected between themselves and also to the rest of the network. If we block the three core nodes, every branch becomes disconnected. This makes the propagation of information inefficient. Regardless of where it starts, it will always be limited to that network branch. This insight will help us later on when trying to understand our results.

As explained before, the clustering coefficient can be correlated with the number of triangles. Because DMS adds a triangle on each step, the resulting coefficient is high.

To help visualise the hierarchy, DMS's triangles can be reduced to a single node. In doing this, there is only a single unique path (without backtracking) between any pair of nodes. See Figure 3.3 for an example. We can maintain the structure of DMS, without the clustering (degree distribution would also change), and it may give us some insight on how each property influences the spread.

From the core, we can quickly reach any node in the network, but if we start in nodes that are further

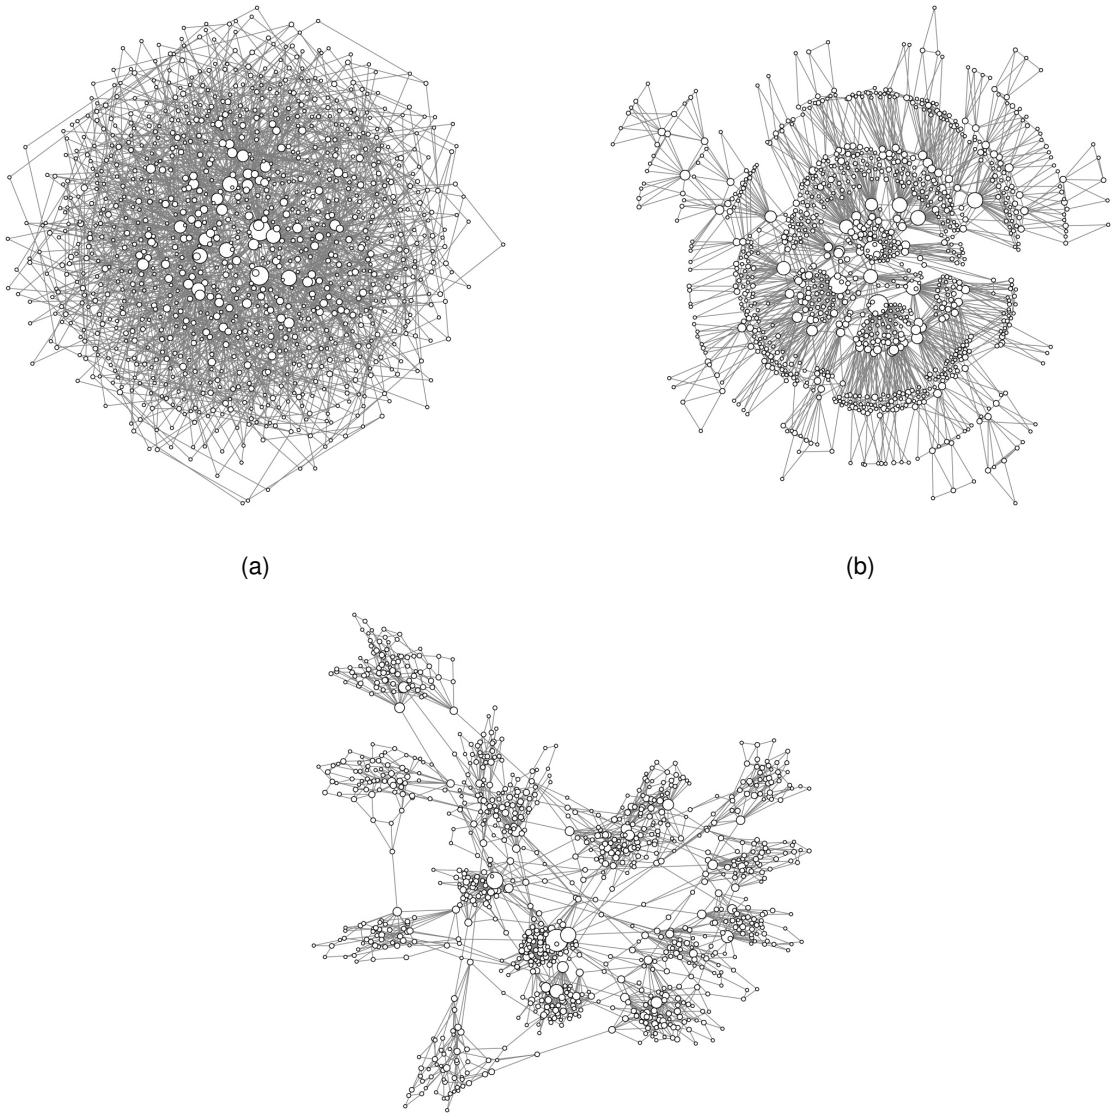

(a) (b)

(c)

Figure 3.2: Example networks with 1,000 nodes for (a) BA, (b) DMS, and (c) LFR. Using Kamada-Kawai force-directed layout [18] (BA and DMS) and Spring layout [10] (LFR). Spring was used for LFR since it shows the community structure more clearly.

away from the core, like leaves, then reaching most other nodes in the network requires having to pass through the core. This gives us an increased Average Path Length by about 1 step compared to BA. It does not increase by much, since most branches aren't that deep, due to the random edge selection having a higher chance of picking edges near the core (preferential attachment). This means that the density of nodes is higher near the core, and makes the APL not increase by much. We can see this by measuring the average of the shortest distance from every non-core node to one of the core nodes, which we can call Average Shortest Distance to Core (ASDC):

$$ASDC = \frac{\sum_{d}^{A} shortest\_path(d, C)}{N - 3} \tag{3.1}$$

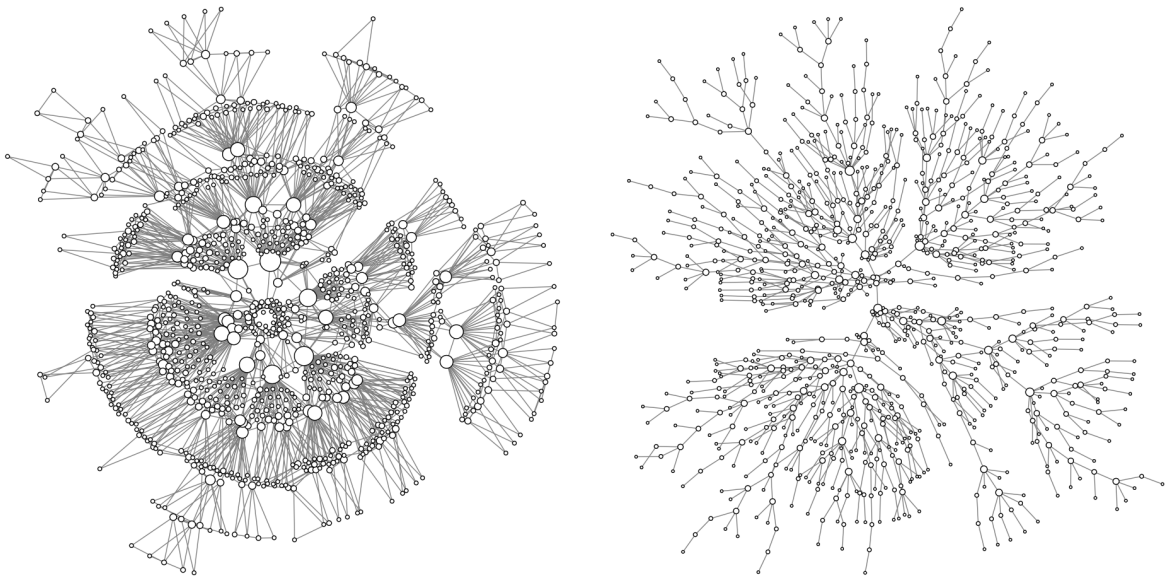

Figure 3.3: Reduction of a DMS network with 1,000 nodes.

where A is the set of all nodes, C is the set of core nodes ($C \subset A$), $d \in A \wedge d \notin C$, and $N - 3$ represents the number of non-core nodes.

We applied this metric to 1,000 random DMS networks and got an average distance of 2.52 steps (standard deviation of 0.19). It's about half of the average path length, and shows just how close most nodes are to the core.

3.3.3 Lancichinetti-Fortunato-Radicchi (LFR)

LFR, similarly to DMS, gives us a different topology from BA and allows us to study how communities can influence the spread of information, yet keeping a power law degree distribution. The LFR algorithm does not necessarily create scale free networks, but if we supply the degree distribution exponent parameter with value 3, it will have the same expected degree distribution for BA and DMS. From this, the other parameters can be picked more freely according to what we want [1]. We chose values that produced networks with communities sufficiently large but not too connected to each other. This allows us to study what effects this community isolation produces, like the polarisation found in "echo chambers" (social groups that repeat the opinion they believe and ignore information that would challenge it). Values chosen besides the defaults of the networkx algorithm were: community size distribution exponent with 2; fraction of inter-community edges at 0.07 (intra-community edges $(1 - mu)\deg(u)$); average degree with 5; and minimum community size at 50 nodes. This gives networks like the one seen in Figure 3.2c. With these values, the algorithm has a propensity to generate networks that are not connected (communities or nodes not connected to the rest of the network). In order to combat this we generated a large number of networks and verified that they were indeed connected. We did this until we had 100 networks saved[2] which were reused throughout the rest of this paper to obtain the results, with the

[1] You can find networkx's LFR algorithm documentation here.
[2] The saved LFR networks can be found here.

exception of the networks of size 10,000. For these larger networks, we tried to generate them using the same parameter values in order to get comparable results. This resulted in a single connected network out of thousands of tries. Given this we decided to save 100 networks whether they were connected or not and the results show that despite this, the values are rather similar as we we will see in the beginning of Sections 4 and 5. Networks with minimum community size of 100, and connected, were also tried but had higher differences in their values, with the highest being over 10%. A community size of 500 was also tried, but the algorithm struggled to generate them. Despite not being connected, we can measure how disconnected the networks are through the use of connected components. Without considering the largest component, on average there were ~9.32 components ($standard\ deviation(s) = \sim 3.72$) which made up on average ~26.5 nodes ($s = \sim 33.29$). Since these other components make up on average ~0.265% of each network, we can assume they will have negligible effects on the results.

For clustering, we can see in Table 3.2 that it is higher than BA but still far from DMS. This increase is assumed to be from the edges inside the communities, since neighbours inside the same community have a higher chance of knowing each other.

The APL (see Table 3.2) on the other hand increases by almost 3 steps, compared to BA. In our LFR we don't have a core like DMS or an abundance of paths between nodes. Instead, we have multiple dense communities, connected by few paths, making the distance to travel between two nodes, on average, that much higher, and the nodes become much more disconnected to others in other communities.

With regards to hubs, the chance of a hub being the connection between communities is higher since having a higher degree means a higher chance to be the one with the inter-community edges, given that we use a very low value for this fraction.

4 Results: External Intervention

If we were to fight misinformation on a population that does not share truthful information and can only learn it, how could we approach this? By learning truthful information an individual would be considered vaccinated. In the classical sense of a vaccination protocol, individuals learn truthful information via a top-down approach. And what if this population was arranged in hierarchical or community fashion, how would that change the outcomes and our options to fight misinformation? This is what we look into in this chapter by using classical vaccination as a means to study how a non-spreading truthful information (the vaccine) can influence the spread of "associated" misinformation.

For the following results we used networks with 1,000 nodes, and for SIRV, we used stifling rate γ and forgetting rate δ of 0.1, and vaccination rate θ of 0 to simulate classical vaccination. The other parameters, infection rate β and vaccination stifling rate ζ, varied from 0 to 0.9, with a step of 0.1. For γ and δ, the value was not changed because they're not the focus of this study. Their value was chosen to account for what they represent, non-vaccine stifling and forgetting respectively. Misinformation always starts in a single node, which is chosen randomly, unless otherwise specified. For each combination of parameter values, 1,000 simulations were executed for statistical purposes, and the results were obtained as a mean over these simulations. For both BA and DMS, we generated 100 random networks at the start of every simulation batch (which usually contains enough simulations to produce a figure) and use those for that batch specifically. As said in 3.3.3, LFR uses the same 100 networks for all results, with the exception of when larger LFR networks are needed for comparison of data for different sizes.

4.1 Network Size

In order to study the effects of network size on the results, we consider the differences between 1,000 and 10,000 nodes. By comparing results for Total Misinformed (TM), the number of stiflers (R) at the end of a simulation, with β and vaccination fraction, we expect that due to the scale-free nature of the three networks, the results should be consistent.

For BA, seen in Figure 4.1a, the lines for both sizes are indistinguishable, which can be attributed to its lack of clustering.

For DMS, seen in Figure 4.1b, the lines show more variation, with 10,000 nodes being always lower. With zero vaccination, the lines are closest, but this is lost on the other vaccination values due to topology. According to the shortest distance to the core nodes, nodes can be placed in ring like structures, with the core being the highest ring. Each ring corresponds to the distance and is made up of two layers (except the core which is a single layer). The internal layer are nodes that were attached to an edge of both nodes belonging to the higher ring. The external layer is composed of nodes where one of the

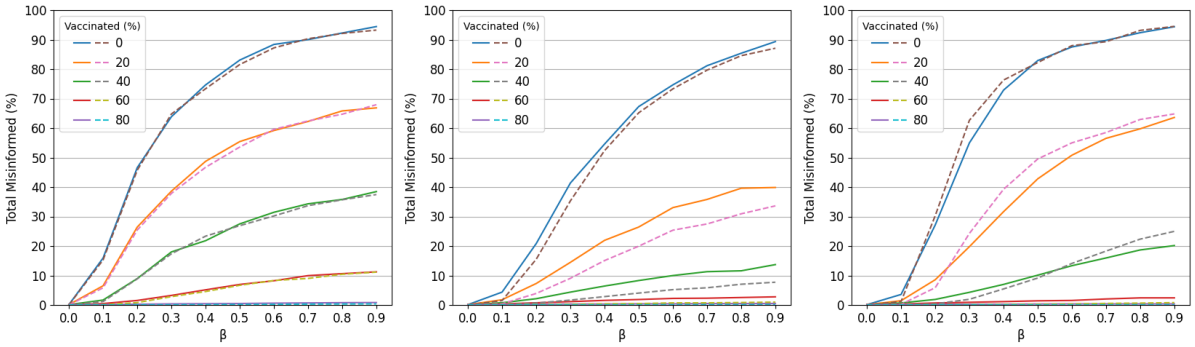

(a) (b) (c)

Figure 4.1: Comparison between results for various networks and sizes for them at 1,000 (solid lines) and 10,000 (dashed lines) nodes. The networks are as follows, (a) BA, (b) DMS, and (c) LFR. Both vaccination rate θ and vaccination stifling rate ζ are set to 0, stifling rate γ and forgetting rate δ are set to 0.1, while infection rate β varies from 0 to 0.9 with a step of 0.1. The vaccination fractions tested are 0.5, 1, 2, and 4. The LFR networks of 10,000 nodes are not connected while the 1,000 sized ones are. For more information regarding this, see Section 3.3.3.

nodes is of the same ring and the other of the higher ring. A depiction of this ring structure can be seen in Figure 4.2. Averaging the number of nodes per shortest distance to the core, we get the distribution of nodes per ring and per layer, as seen in Figure 4.3a for 1,000 and Figure 4.3b for 10,000 nodes. For the former, the second ring has the most nodes, while on the latter, the third ring is larger. As we increase the size further (100,000 and 1,000,000) the largest ring also increases. With 10,000, more nodes are in rings up to, and including, the largest ring, when compared to 1,000. This leads to more nodes being vaccinated in the lower rings, causing more disruption and blockage of misinformation spread to the higher rings. This can be attributed as the cause for the lower TM for 10,000 when vaccination is applied. The blocking of branches only happens when both nodes of an edge are vaccinated. This is rare with random vaccination but becomes more commonplace as vaccination fraction increases. However, just vaccinating a single node in a branch can disrupt the misinformation spread, because it needs to go around. This will slow down the spread momentum, given that infection is not 100% and there's stifling and forgetting. In spite of this, the trends of the lines are always the same. Figure 4.3a is also way of visualizing our $ASDC = 2.52$ for 1000 nodes since most nodes are close to the core.

For LFR, seen in Figure 4.1c, the lines follow the same trends, but there is an apparent relation between both sizes. The lines for 10,000 start lower on the low β's and past a, sort of, transition point (which varies depending on the vaccination fraction) becomes higher which maintains until 0.9. Exceptions for this are for vaccination 0 where after becoming higher, the lines then match until 0.9, and in the case of vaccination fraction of 0.6 where the 10,000 line never surpasses the other. For 0.8 the nearly horizontal lines are indistinguishable. We believe this relation comes from the way the 10,000 sized networks were generated, as explained in 3.3. The minimum community size was kept at 50 in order to be able to generate the larger networks, and in doing so did not accompany the 10 fold increase

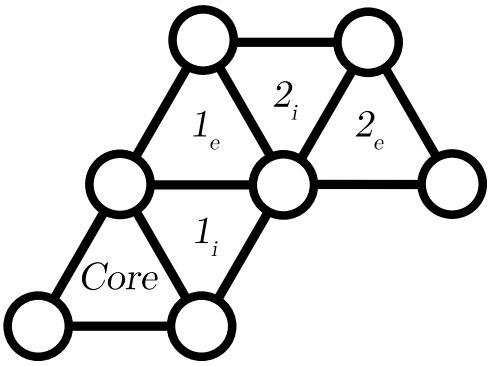

Figure 4.2: Simple depiction of the ring structure in DMS. Every node not belonging to the core belongs to K_x where K is the ring (shortest distance to the core) and x is the internal (i) or external (e) layer. The layer is determined by the number of edges to the lower ring, where 2 means internal and 1 means external.

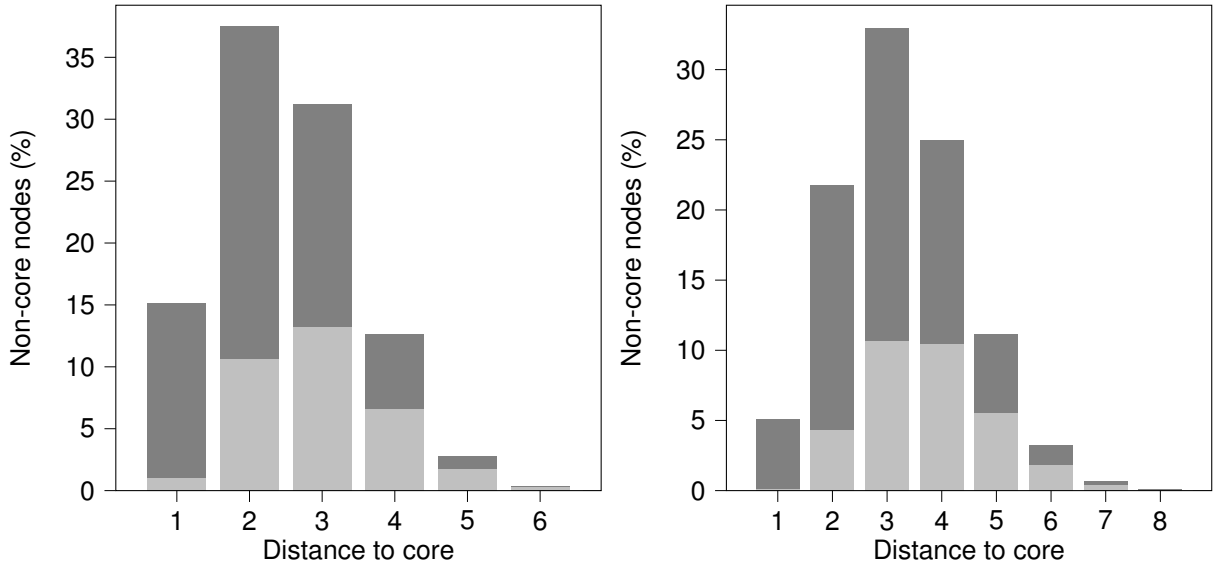

Figure 4.3: Distribution of non-core nodes for the distance to the core in DMS networks with (a) 1,000 nodes and (b) 10,000 nodes. Each bin is divided into the percentage of nodes belonging to the internal layer (light grey) and external layer (dark grey). The values were averaged over 1,000 networks each. Distances with less than 1 averaged node were hidden.

in the number of nodes. Since communities are able to be smaller than 500 nodes, and due to the power law distribution for their sizes (with exponent of 2), the number with size lower than 500 will be significant to affect the results. For lower beta the misinformation will spread less as it'll have a greater chance of hitting a smaller community and die, before spreading further. On the higher beta values we get a different effect due to the increased inter-community edges which offer more paths for spread. As we apply and increase vaccination, we shift the transition point further into the higher β's as more infectivity is necessary to make use of the higher inter connectivity. At vaccination 0.6 (and perhaps higher), this point moves to the upper bound of β.

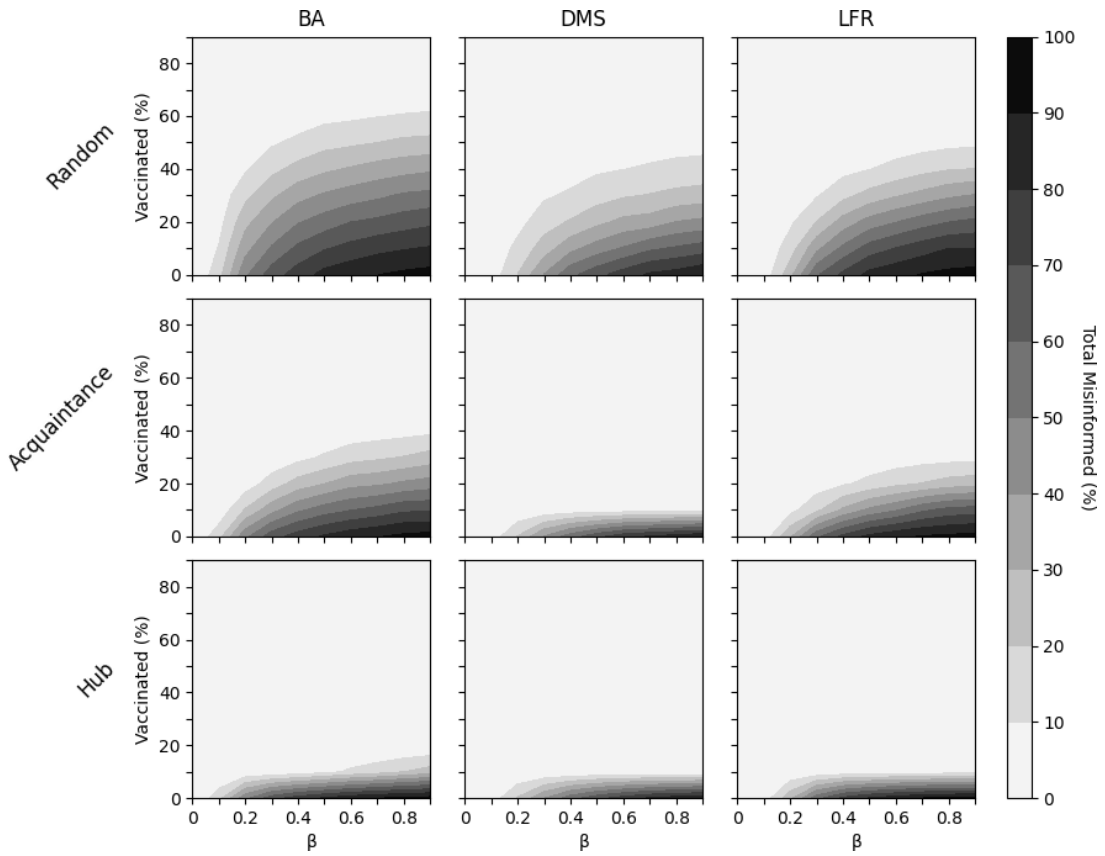

Figure 4.4: Contour graph depicting how different strategies of vaccination (rows: Random, Acquaintance, and Hub) affect different networks (columns: BA, DMS, and LFR) considering vaccination fraction and infection rate β. The colours shown represent values for Total Misinformed, which is the percentage of stiflers at the end of a simulation. 1,000 repeats for each point of data were done. Vaccination fraction and β vary from 0 to 0.9 with a step of 0.1 giving us 100 data points. Vaccination rate θ and vaccination stifling rate ζ are set to 0, while stifling rate γ and forgetting rate δ are set to 0.1.

Given this, the size for the networks was chosen to be 1,000 nodes, because it provided faster results (faster network generation and simulation) while also providing us with relevant information for the relationships between the variables of interest.

4.2 Vaccination Fraction

For fighting misinformation in a classical vaccination scenario, we have to employ strategies on what nodes to vaccinate and how many. We'll be focusing on Random, Acquaintance and Hub strategies. Random works as a baseline for comparison. By varying the infection rate β, we obtained results regarding all 3 strategies and all 3 networks as seen in Figure 4.4. The data points have intervals of 0.1 for both axis, which means intervals with large changes may not be indicative of the actual tendencies, as seen in the Hub data. This will be addressed further on.

We can see that, as β increases, the gains from increasing vaccination become minimal past a certain cutoff point. This point depends on the network type and the strategy employed. Above 0.50 β

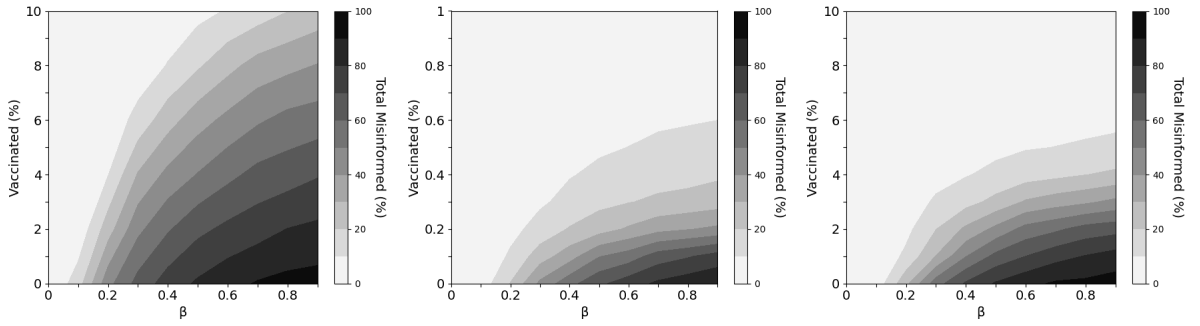

Figure 4.5: Contour plots depicting how the Hub strategy of vaccination affects the various networks (a) BA, (b) DMS, (c) LFR, given vaccination fraction and infection rate β. The vaccination fraction scale for BA and LFR (0 to 0.10) is different from DMS (0 to 0.010) to show the contrast between them. The colours shown represent values for Total Misinformed, which is the percentage of stiflers at the end of a simulation. For the three networks, β varies from 0 to 0.9 with a step of 0.1. Vaccination fraction has a step of 0.01 (BA/LFR) and 0.001 (DMS). Vaccination rate θ and vaccination stifling rate ζ are set to 0, while stifling rate γ and forgetting rate δ are set to 0.1.

it appears that increasing the vaccination fraction yields minimal gains, independent of the network or strategy.

Random Strategy. As expected, this strategy works rather poorly on the various networks, since it doesn't try to take advantage of the underlying structure like Acquaintance or Hub. Nevertheless it behaves better in DMS and LFR than in BA. In DMS, despite the cutoff value being similar to LFR, the overall TM has a slower increase as the vaccination fraction decreases, which as discussed below, seems to be due to its hierarchical structure.

Acquaintance Strategy. With this strategy, since it has a higher chance of vaccinating hubs, the cutoff points drop considerably for all network types. While the chance of vaccinating hubs is relatively the same across all of them, since they have similar degree distributions, DMS has a significant drop when compared to the others. This difference can again be attributed to the network's topology.

Hub Strategy. For this strategy, where vaccination blocks the main highway nodes for information, we get another considerable drop. All networks now generally require less than 10% of the population to be vaccinated to effectively curb the misinformation spread. Since it falls between the interval of data points, it becomes impossible to see the actual data tendency, but if we redo results for this interval at smaller steps, we can see a clear difference between the three (Figure 4.5, with a difference in vaccination fraction scale between DMS and BA\LFR to show the contrast between them). We can notice the different orders of magnitude between DMS and BA\LFR, where DMS goes from 0 to 1% and BA\LFR goes from 0 to 10%. While the Hub strategy on BA is still not very effective at these lower fractions, DMS and LFR have a clear difference in effectiveness, with the former showing a middle cutoff point between the higher and lower valued areas at 0.2 to 0.3% (2 to 3 nodes, respectively), while the latter shows a cutoff point at about 5%.

4.2.1 Structure

How does the structure influence these results? With DMS, we know the extreme importance of the core nodes in propagating information, so any strategy that can vaccinate hubs with a higher chance can take advantage of this. The core nodes will most likely be the largest hubs, and it explains the striking difference in results between Random and Acquaintance\Hub strategies. As said in Section 3.3.2, by vaccinating 2 or 3 core nodes, we can cutoff a great number of branches from the rest of the network, or even every other branch, respectively. By doing this, the misinformation is contained inside those branches and is effectively stopped from spreading through the network. This is what we see in Figure 4.5b since for 0.2% (2 nodes) to 0.3% (3 nodes), we have a clear cutoff point for the high TM values. The likelihood that those nodes are in the core is very high. Furthermore if we keep increasing the vaccination fraction for the Hub strategy we will vaccinate even more hubs, which will be the entry points into multiple sub-branches, and those will become disconnected too, further curbing the misinformation spread. For the Acquaintance strategy the same applies, to a lesser extent, since the possibility of vaccinating hubs is still higher than Random, but not near as high as Hub. While Random behaves poorly when compared to these strategies, it still behaves better than in BA and LFR because of the structure. The likelihood that a random node will block a branch or sub-branch is non-negligible due to most of nodes being located in the middle of the hierarchy (second and third ring as seen in Figure 4.3a). Since there are no other paths into those blocked areas, on average, it's still more effective comparably.

In BA on the other hand, given the higher abundance of paths between nodes in the network, randomly vaccinating nodes is not very effective, since even if we do manage to block a hub node, the misinformation can still spread around that blockade. Thus, the vaccination fraction has to be higher to compensate for this in order to block more paths. With Acquaintance, we can more likely block information highways so the vaccination fraction needed overall drops, but is still high comparably. With the Hub Strategy it decreases again, considerably, although we still need at least around 10% vaccinated hubs (for high β's) to reach low TM values.

With LFR we discussed before that hubs have a higher chance to be the connections between communities. Given its community topology, with few paths connecting communities, the higher the chance of vaccinating a hub the lower the TM, like with BA and DMS. With a random vaccination strategy, a high enough fraction will tend to block paths, but since there are less paths than BA, we will need less vaccinated nodes overall. However, as we lower the vaccination fraction the results are similar to BA, as paths are less likely to be blocked so we reach similar TM values. Further on, in Chapter 5 we will see that, despite the propagation being different over time, the end results are similar. With Acquaintance, we see identical results for low vaccination fractions, again, but the cutoff point is lower, on account of a higher probability of blocked paths for higher fractions due to vaccinated hubs. With Hub we can use the same reasoning. When we vaccinate hubs, the probability of paths being blocked is higher while the cutoff point is still around 50 nodes (see Figure 4.5c). This means there are hubs that have no inter-community connections, despite being likelier. It follows that these nodes are only good for curbing the propagation inside its community once misinformation reaches it. Thus, this requires more hubs to

be vaccinated.

Clustering Coefficient

It's difficult to separate clustering from structure. As clustering increases so does the number of nodes with similar states that interact with each other. The increase in clustering comes with a higher polarisation, which we will discuss in Chapter 5. In this case, since vaccination does not spread, misinformation tends to takeover more clustered areas, forming a polarised zone. This will, in turn, increase the stifling that happens, because spreaders will stifle each other, while also forgetting and starting to stifle those around them. More clustered zones or communities, which are dominated by the spreaders, will then begin to die out from the inside, due to stifling, with a "virus" like behaviour. This should theoretically force misinformation to disappear sooner.

4.2.2 Concluding Remarks

From these results, we can see why the structure is very important. Hierarchical structures have an inherent strong resistance against misinformation spread, while networks with low clustering show a very weak one. The difference between these "extremes" is the degree of order that they have. We assume the correlation is that: the more order that there is in the structure of a network, the more resistant it becomes to the spread of misinformation, and the easier it becomes to vaccinate against it.

Regarding the separation of structure and clustering, some rough tests where applied to both BA and DMS, by counting the number of stifling interactions that happened on average. The values where fairly similar, which could mean that clustering does not affect amount of stifling. Regarding DMS specifically, on Section 3.3.2, triangle reduction was briefly discussed. This could, in theory, remove the clustering from the DMS networks while leaving the branch like structure intact. From this we could generate sufficiently sized networks so that after reduction they would have a similar number of nodes to non-reduced DMS networks that we can compare results to. We would be essentially separating these two connected parts of the network and it might give us some insight into how the structure affects the spread with and without clustering. Both these points could be the focus for future work, and will not be addressed here.

4.3 Vaccination Stifling

So far we modelled vaccinated individuals as uncontactable nodes to spreaders, but we can instead let them be contacted in order to stifle spreaders into changing their behaviour. We can conceive from this contact, some sort of social pressure, maybe through dialogue or exposure to a different opinion, the misinformed node will change their attitude or even opinion and turn into a stifler, no longer believing that what they believe should be spread. This will then cause that node to stifle its spreader neighbours, and so on.

In order to get results, we fixed the infection rate β at 0.8 and varied the vaccination stifling rate ζ and vaccination fraction. β was fixed at a high value in order get a large misinformation spread, in order

to have a higher TM baseline (no stifling, $\zeta = 0$) to compare to when stifling is introduced. This stems from the idea that the more misinformation spreads, the more it will come in contact with vaccinated individuals and the more stifling events will occur, which possibly lead to higher decreases in TM. This becomes more apparent when we look at stifling on Section 5.3. The results can be seen in Figure 4.6. We can see some similarities to Figure 4.4, like DMS having the lowest overall TM, followed by LFR and then BA. Another is the sharp drop from Random to Acquaintance for DMS.

With DMS, starting on random, we have a clear difference in the slope of the areas. The sharp drop in TM on low ζ's is due to clustering and how nodes are spread out via random vaccination. Since vaccinated nodes are spread out randomly on the network, they are linked to a higher quantity of nodes because of clustering. We can call this amount of edges the area of contact between vaccinated and ignorant/spreaders. As vaccination fraction increases, this area will increase until the point where most neighbours of vaccinated nodes are other vaccinated, where the increase will stagnate. But even with a low vaccination fraction, this area will be large due to the high probability of choosing nodes in the second and third ring, which have a tendency for high degrees but still lower than the first ring (which contains most of the hubs).

As vaccination fraction increases, so does the chance that there will be a stifling event early on in the misinformation life, which will affect its spread considerably, unlike a later event. This is again due to large area of contact that random vaccination produces with the high clustering. Therefore, even with a low ζ, the clustering forces stifling events to happen more often and, thus, will have a greater effect. If a single vaccinated node is at the entrance of a branch, then misinformation will have to spread around it via the other entrance node. As it goes around it, it increases the area of contact between misinformation and vaccination even further, leading to more events. As ζ increases, the events are more successful which means they will stifle misinformation much sooner, but the area of contact will not grow as a consequence. The reason why the decrease in the vaccination fraction for each TM area stagnates is because when it becomes lower, vaccinated nodes can keep the misinformation from spreading easily (via the high ζ) through the nodes around them, but can't stifle the areas where there are no vaccinated nodes. As the fraction increases, the areas untouched by vaccination will disappear and it can keep the misinformation from spreading at all.

The effect of clustering on these results explains why we don't see this in the other two networks, and might explain why LFR has a slightly bigger slope than BA.

With Acquaintance, given its high probability of vaccinating hubs, the effects of stifling become negligible due to the structural advantage in stopping the spread. With Hub we see the same to a higher degree.

BA and LFR give similar results in terms of the data tendency as ζ increases. The reason why we do not see a sharp initial drop as in DMS is, again, due to the structure. In BA there are more paths to use to spread around the network and misinformation is not forced to spread through the neighbours of vaccinated in order to reach certain nodes. Even if they do come in contact with a vaccinated, the next node to be misinformed will, most likely, not be a neighbour of that vaccinated node - with vaccinated hubs the chance is higher, but still very small due to the almost null clustering coefficient of BA. Therefore

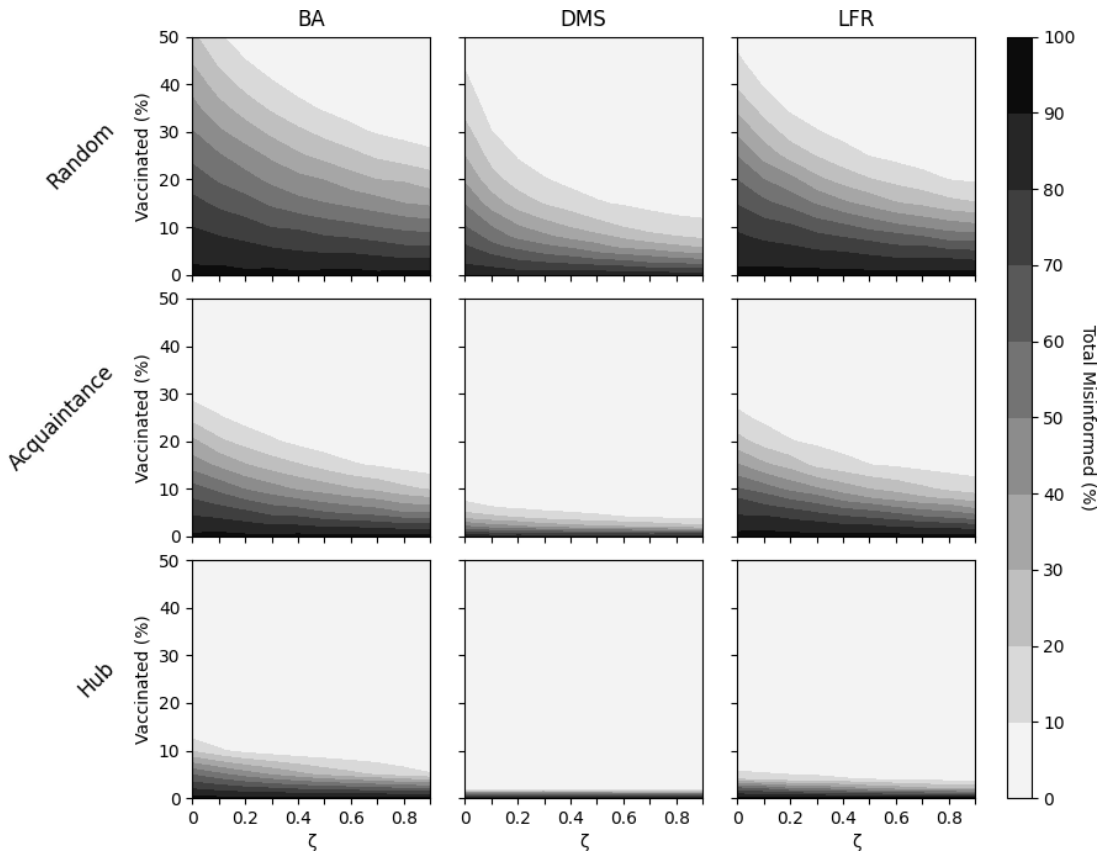

Figure 4.6: Contour graph depicting how different strategies of vaccination (rows: Random, Acquaintance, and Hub) affect different networks (columns: BA, DMS, and LFR) considering vaccination fraction and vaccination stifling rate ζ. The colours shown represent values for Total Misinformed, which is the percentage of stiflers at the end of a simulation. With infection rate β fixed at 0.8, ζ varies from 0 to 0.9 with a step of 0.1, while vaccination fraction varies from 0 to 0.1 with a step of 0.02 and 0.1 to 0.5 with a step of 0.05. The difference in steps is to account for granularity between the interval 0 to 0.1, especially in Acquaintance (DMS) and Hub, which helps for comparisons. Each data point was averaged over 1,000 repeats. Vaccination rate θ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1.

there are less chances of stifling and it requires more vaccination percentage, or higher ζ, for stifling to happen successfully and more often. In LFR, the community structure restricts the spreading speed (discussed more in depth on Section 5.2.1) and misinformation tends to become stuck in communities until it gets a chance to "jump" to the next. Stifling will happen later, but not necessarily less than in BA, since vaccinated nodes will be spread out in the various communities (especially with the Random strategy). Given that stifling happens later on, it is less useful for reducing TM. In the case of the Hub strategy for LFR, as hubs have a lot of connections inside the same community and are most likely the entrances to it, when misinformation enters a community it will be in immediate contact with the vaccinated node and also, as it spreads, it will maintain contact with it due to its high degree. The higher clustering, compared to BA, will also help in causing more stifling events.

4.3.1 Concluding Remarks

From these results on vaccination stifling with classical vaccination we gather that whether or not stifling has an impact depends, mostly, on the structure, but also on the virulence of the misinformation. Misinformation that does not spread much will have less stifling events and thus less effect on the TM that is reached. But structure (and clustering coefficient) plays an even more important role depending on how it affects spreading. As we saw with DMS, since it funnels misinformation into having prolonged contact with vaccinated individuals it is very useful for reducing TM, when used with a Random strategy. But when we use Acquaintance or Hub, the effect of structure in these strategies causes stifling to have negligible effect. In BA and LFR, the effect of stifling is much more reduced due to not having this prolonged contact effect. In the case of LFR, this effect can happen with the Hub strategy since hubs are connected to most of the community and there is more clustering.

5 Results: Self-Organized Intervention

What if, instead, individuals in the population tried to spread the truthful information in order to try to reduce the effect of misinformation? What if they tried not only to convert the ignorants to the truth, but try to stifle misinformation spreaders into stopping their spread? This is what we seek to study in this chapter with spreading vaccination and also stifling vaccination.

In this chapter we will also use the concept of timescales, which allows us to give different speeds to the two types of news. If we want to study what happens when vaccination spreads, at two times or four times the speed of misinformation, we can use this concept.

For this chapter, for the results obtained, we kept the same size of networks and values for stifling rate γ and forgetting rate δ in SIRV, as in Chapter 4. For stifling, vaccination stifling rate ζ is kept at 0 until we discuss its results specifically. Both misinformation and the vaccine start in a single node, since we're simulating the bottom up approach of vaccination. The networks used also remain the same, with BA and DMS having 100 networks generated randomly per simulation batch and using the same 100 LFR networks.

5.1 Network Size

We verify again for any effect from the size of the networks by comparing with results for 10,000 nodes, which can be seen in Figure 5.1. In BA we start to see some bigger differences in timescales 2 and 4, but still close enough and with the same tendency to not be significant. For DMS, the lines are very similar so there is no significant difference there. In the case of LFR, the differences are not significant, but we see the same tendency for networks of 10,000 nodes, of lower values for lower β's and higher values for higher β's, as seen in Section 4, for timescales 0.5 and 1. The same doesn't happen for timescales 2 and 4, so the initial reason assumed for this behaviour in Section 4.1 might not be correct, and requires further investigation in future work.

5.2 Timescales

In order to vary the speed of vaccination against misinformation, we can implement timescales through the use of random number generation within an interval. The interval we use is $[0, 1 + w]$, with w being the timescale. If the number is in $[0, 1]$ then we do a batch of misinformation events (transitions). If, instead, it's in $]1, 1 + w]$ then we do a batch of vaccination events. If w has value 1, each

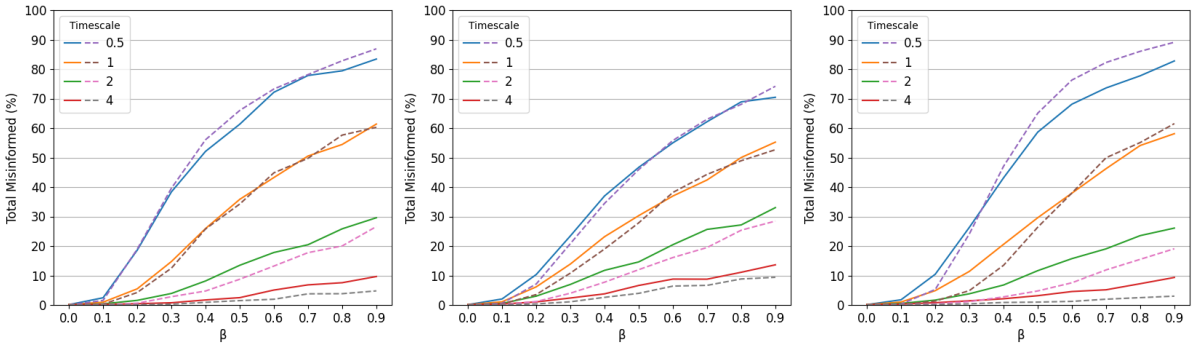

Figure 5.1: Comparison between results for various networks and sizes for them at 1,000 (solid lines) and 10,000 (dashed lines) nodes. The networks are as follows, (a) BA, (b) DMS, and (c) LFR. Vaccination rate θ is fixed at 0.5, while infection rate β varies from 0 to 0.9 with a step of 0.1. The timescales tested are 0.5, 1, 2, and 4. The LFR networks of 10,000 nodes are not connected while the 1,000 sized ones are. For more information regarding this, see Section 3.3.3. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1.

information spreads at the same speed (on average). With value 0.5, misinformation spreads two times faster. With 2, the reverse happens. With this, we can try to understand how individuals, less or more active on the side of vaccination, can affect how misinformation spreads.

Varying timescale can also be seen as an indirect way of varying the vaccination rate θ. By increasing timescale from 1 to 2, we could in theory increase θ from 0.5 to 1 to have the same effect. But for a θ higher than 0.5, it would always lead to an increase to 1, since the probability of a transition cannot surpass that. This means that timescale is, in fact, more than just a fancy way of varying θ. Increasing timescale for the same θ means that vaccination events, which depend on it, will happen more. With a θ of 0.6 (with timescale 2), if the transition fails then we will more likely try again sooner, before misinformation can spread. If it succeeds, then on the next step (likely vaccination) it will have the possibility of spreading even further. So not only can transitions fail but vaccination can spread more. If we were to instead double θ then the events would always happen regardless, but not necessarily spread faster (in terms of the upper ceiling with timescales).

Instead of varying just theta, changing timescales could also be modelled as just variations of all the probabilities. As discussed above, since increasing timescale from 1 to 2 for θ of 0.5 is not the same as increasing θ to 1.0, it could in fact be that we need to increase θ to some value between 0.5 and 1.0 and also reduce β's value, to get the same results. It follows that the other probabilities might also need to be changed. What exactly those changes would be is difficult to know since we're trying to combine two systems (transition probabilities and timescale) into just one complex system. To know how this system varies compared to timescale changes would take considerable study and effort to map and equate it. Due to this, we chose to use the simpler system.

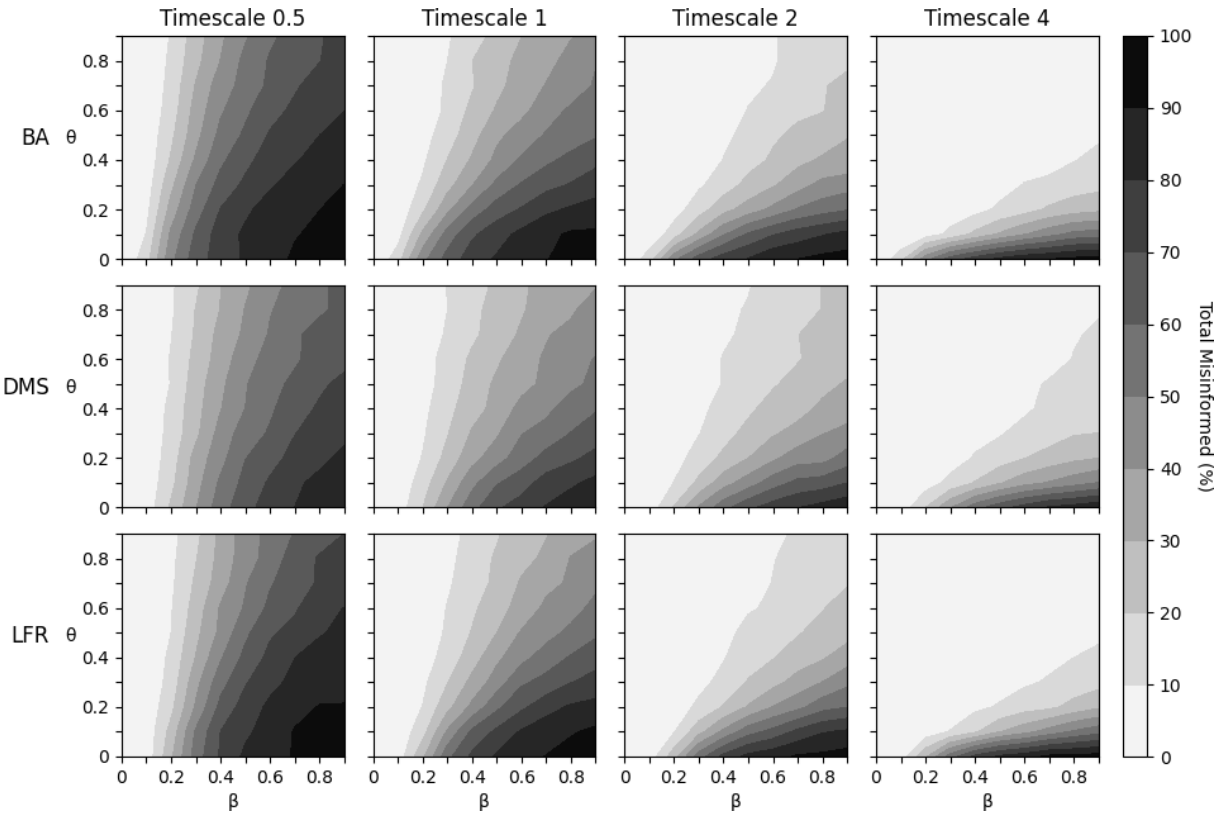

Figure 5.2: Contour graphs depicting Total Misinformed values for various networks (rows: BA, DMS, and LFR) and various timescales (columns: 0.5, 1, 2, and 4) given infection rate β and vaccination rate θ, with both misinformation and vaccination starting in a random node. Both β and θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1.

5.2.1 Random Misinformation and Random Vaccination

Considering different timescales for the spread of vaccination through the use of the timescale factor, how does it affect misinformation spread? By varying timescale for values 0.5, 1, 2 and 4, and vaccination rate θ and infection rate β from 0 to 0.9, with a step of 0.1, we obtained the results which can be seen in Figure 5.2. Overall, independent of the network, by increasing timescale, we increase the effect that θ has. Thus, for lower values of θ we can reach the same TM values by increasing timescale. On lower timescales and lower β, increasing θ has little effect on TM. As β increases, so does the effect that θ has on lowering TM. Nonetheless, increasing timescale provides the largest decreases in TM, regardless of θ or β, since individuals will "work" faster to spread the truthful information.

Similar to what was seen in Chapter 4, DMS has lower values overall compared to BA and LFR, but as timescale increases, the increased effect to vaccination is lower. At timescale 2, the three networks behave similarly (with the exception of very low θ), but at timescale 4, DMS behaves worse for higher β. This can be explained by the structure and will be discussed below. Another clear result is that BA and LFR behave very similar with regards to TM, but given their different topologies, the way the misinformation/vaccination spread on the networks must be fundamentally different. This will also be discussed in the following section.

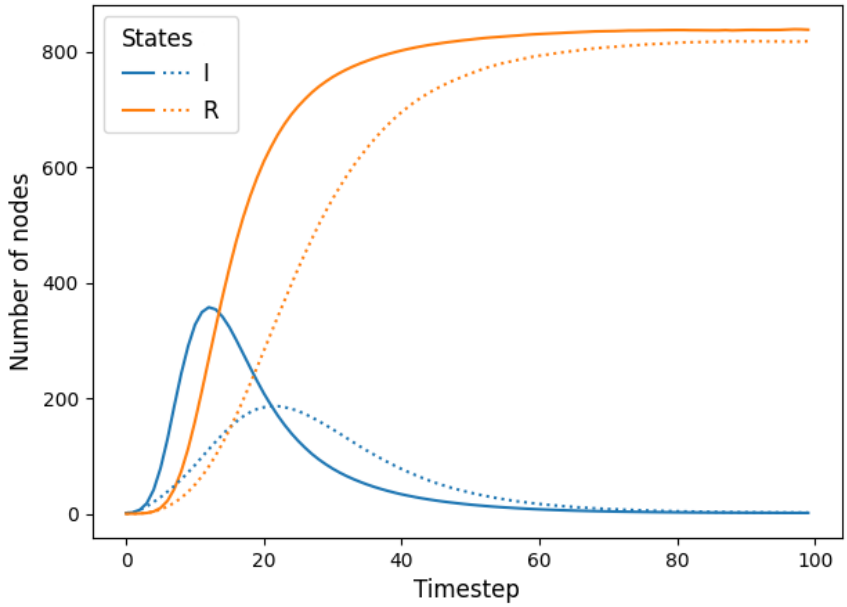

Figure 5.3: Progression over time of the number of nodes per state (spreaders and stiflers) with random vaccination and misinformation, where both start with a single node. The solid lines represents BA states and the dotted lines represent LFR. The data was obtained with infection rate $\beta = 0.5$ and vaccination rate $\theta = 0.2$, stifling rate $\gamma = 0.1$, forgetting rate $\delta = 0.1$, vaccination stifling rate $\zeta = 0$, averaged over 10,000 repeats.

Structure

Starting with DMS, it gives lower TM values overall when compared to BA and LFR for low timescales, since the hierarchical topology makes the spread of information more difficult when nodes can be blocked (vaccinated). As we saw in classical vaccination, BA gave us worse results because it had more paths for the misinformation to spread through and it required a higher vaccination fraction to achieve similar values to DMS. The same happens here, as misinformation cannot spread very effectively, especially now with the vaccination spreading by itself. Nonetheless, given the structure and that both informations start with 1 node, whichever information reaches and controls the core first will tend to control the rest of the network, since it will cutoff the opposite information from spreading outside of its branch or sub-branch. As timescale increases, what helped DMS to lower the spread of the misinformation, now hinders the spread of the vaccination, leading to the higher values (compared to BA and LFR) on timescale 4, and slightly for timescale 2. Despite spreading, on average, at 4 times the speed of the misinformation and most likely reaching the core first, the vaccination cannot stop the misinformation from spreading on the branch or sub-branches that misinformation has already blocked, since there are no paths around that blockage, like in BA and LFR. On higher β, the misinformation can spread more effectively despite being slower and gain more ground before vaccination can stop its advance.

While BA and LFR give similar results, the propagation of information is fundamentally different,

because of their structures. In BA the information can spread more freely and faster while in LFR it is forced into the few inter-community paths in order to spread between communities. This is what we see if we look at the progression over time of nodes per state, specifically I and R, as seen in Figure 5.3. The peak of spreaders is reached much sooner in BA and, naturally, a higher value because of this. The propagation in LFR is much slower, because the spreading between communities takes longer, but also because spreading inside of a community as a random starting node (most likely a node with low degree due to the distribution) takes more time. Even despite this difference in propagation, the TM (number of stiflers at the end of the simulation) remains similar, albeit slightly lower for LFR. Most of this is maintained even with misinformation starting in a hub, which will be discussed further down. It is also worth mentioning that as the number of spreaders grows, so does the number of successful stifling and forgetting events. As more stiflers appear, stifling events will keep increasing and misinformation will die faster.

Clustering Coefficient

As said during classical vaccination, it's difficult to separate clustering and structure, and thus difficult to know how it affects our results. The triangle reduction could provide insight into it, but only for DMS specifically. For LFR, when we look at polarisation in Section 5.2.4, both clustering and structure are the reason why communities become polarised, so it's another example of the difficulty in separating the two.

5.2.2 Hub Misinformation

What if, however, the starting node for misinformation is someone important/influential like a hub? How could we fight it according to our current knowledge? If we consider that misinformation starts in the largest hub in the network, and vaccination continues to start in a random node, we can repeat the results we obtained for Section 5.2.1 with these changes and we obtain the results seen in Figure 5.4. Overall, misinformation spreads much more and reaches higher TM values. Furthermore, as timescale increases so does the general amount of TM, but to a considerably lesser degree than what we saw with random misinformation. The effect that the vaccination rate θ has, especially on lower infection rates (β) and timescales, is significantly reduced, since by increasing it we get very minimal to negligible returns. DMS continues to give lower results, while BA and LFR give similar ones.

DMS

Starting in the largest hub, misinformation will most likely start in one of the core nodes, otherwise it will start in one close to it. This allows it to have the upper hand in spreading through the hierarchy since it can easily take over the core and parts of branches nearer to the core. This way, it nullifies any chance that the vaccination has of spreading and curbing the spread of misinformation, since it will be contained inside of the branch where it started. For low β, where misinformation doesn't spread much, vaccination doesn't matter, since both sides will not come in contact enough times, on average. Any path that vaccination blocks won't matter much, on average, since misinformation won't spread

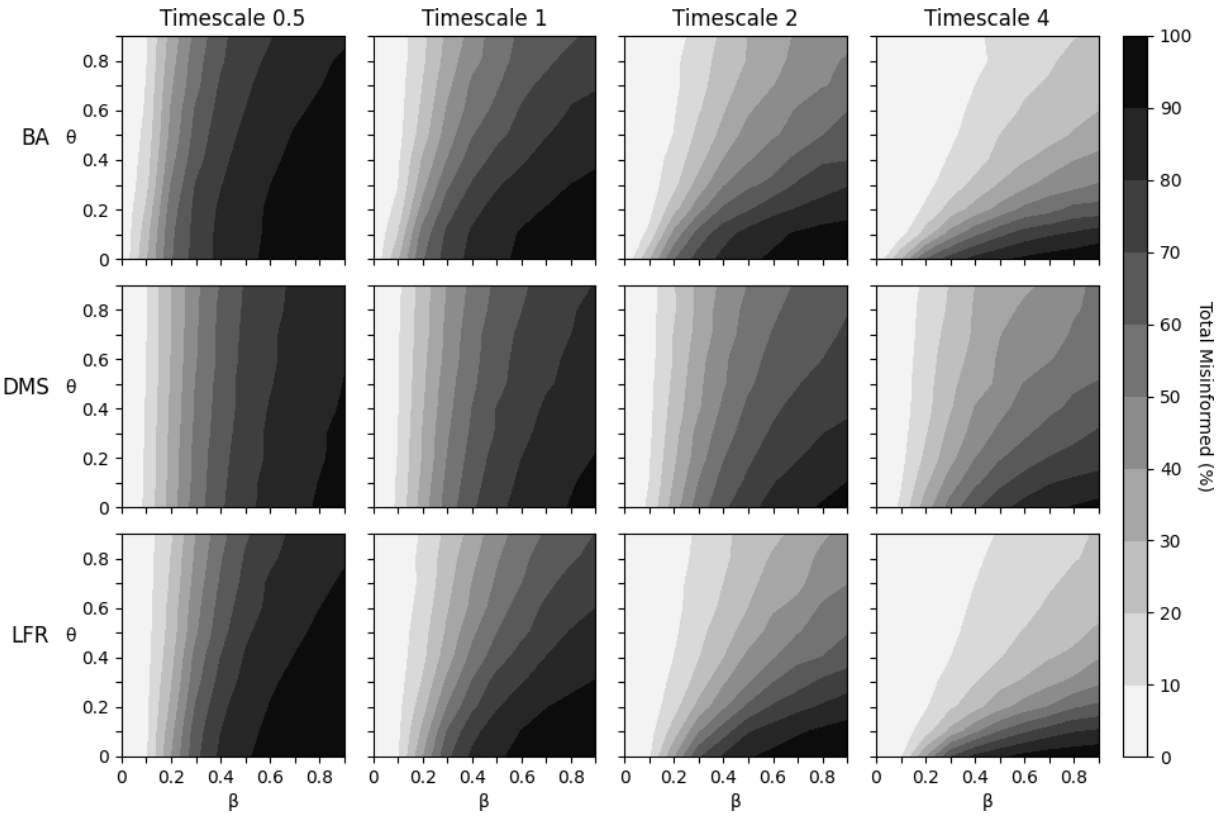

Figure 5.4: Contour graphs for Total Misinformed with different networks (rows: BA, DMS, and LFR) and timescales (columns: 0.5, 1, 2, and 4). Both infection rate β and vaccination rate θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1. Vaccination starts in a random node while misinformation starts in the largest hub.

enough enabling it to reach the zone that was blocked. Since misinformation will also die out quickly, vaccination does not spread much. As we increase β, misinformation will spread more and further from the core and vaccination blockage will actually matter, as seen from the decrease in TM as we increase θ. Furthermore, misinformation can reach beyond those blockades with low vaccination, but as we increase it, branches become blocked and it spreads less.

As timescale increases, the effect of θ will increase too, more notably on the higher β values, since the spreading vaccination will block more and more branches that misinformation could spread to. But, overall, if misinformation starts in the largest hub (core or near the core), increasing either θ or the timescale does not give us large gains in reducing TM. Compared to BA and LFR, and given the importance of the core in DMS, it becomes very difficult to combat misinformation, despite best efforts, unless it is an extremely "infective" misinformation, which we can then reduce its presence slightly. Despite being improbable, it is still possible that with higher timescales and θ, that vaccination might be able to quickly make its way to one of the core nodes, if there is a clear path to it and it starts near the core (remember the average path length to core at around 2.5 steps, and the node distribution per ring), before misinformation can spread to it. Especially if misinformation starts in a hub that is not part of the core. This way, vaccination could block entire branches and sub-branches.

BA and LFR

Again we obtain similar results between these two networks and, as seen before, the way that propagation works is fundamentally different. On lower timescales, increasing θ for low β returns small gains in reducing TM for both networks. On BA, the blockage by vaccination gives the same effect as in DMS, where misinformation doesn't spread far enough for it to matter. Unlike DMS, even if the misinformation is blocked it can still go around it due to its structure. As we increase β, the misinformation starting in a hub can take advantage of the structure to spread faster and more effectively since it can reach a large number of nodes from the starting position. Even as vaccination increases, the structure still helps misinformation resist it due to the possibility of going around it. In LFR, starting in a hub means starting in the focal point of one community, but besides spreading to that whole community quickly and effectively, it is still bottlenecked into the inter-community edges. However this will make it spread slower to other communities, the same applies for vaccination, and since it starts in a random node, it will take longer for it to take over the community where it starts, and to spread outside of it, so misinformation will generally have the upper-hand in spreading.

As timescale increases, the effect of θ increases, and misinformation is overtaken by vaccination. Even though it starts in a hub, with a high enough θ and timescale, we can reduce it significantly on these networks types, since while these networks help in spreading misinformation, they can also help spread vaccination with high enough timescales.

Since they behave similarly, we can again look into how they progress over time, for the number of nodes per state, with a focus on spreaders and stiflers. An example of this can be seen in Figure 5.5. Like before, despite reaching similar values of TM, the spread is much slower in LFR due to the community structure. Furthermore, TM for LFR continues to remain below BA. But there is one fundamental difference between the random and hub misinformation, and it is how LFR starts. Before, with random vaccination, the line for spreaders for LFR started in a different trend than BA. Starting in Hub, for a few initial timesteps, the lines match, until the same tendency for LFR (as seen in random) begins. What this tells us is that the spread from the initial hub is identical for both networks. What this might mean is that spreading to the entire starting community in LFR might be the same as spreading from the hub in BA. But without a further look into when exactly this point of change (same growth to slower growth) happens, it might be that this initial spread is most likely due to the 1st order neighbours of the largest hub (which will on average have the same degree for both networks given the degree distribution), and after that point it slows down. But this also seems counter intuitive that it would slow down after spreading to so many nodes. Perhaps having a higher clustering coefficient in the communities slows down the spread, since spreading to more nodes in the same community will not give us more new untouched nodes like it happens in BA.

5.2.3 Hub Vaccination

Like misinformation, the truthful information can also start in the highest hub. We can return misinformation back to starting in a random node, and repeat the results again to get what can be seen

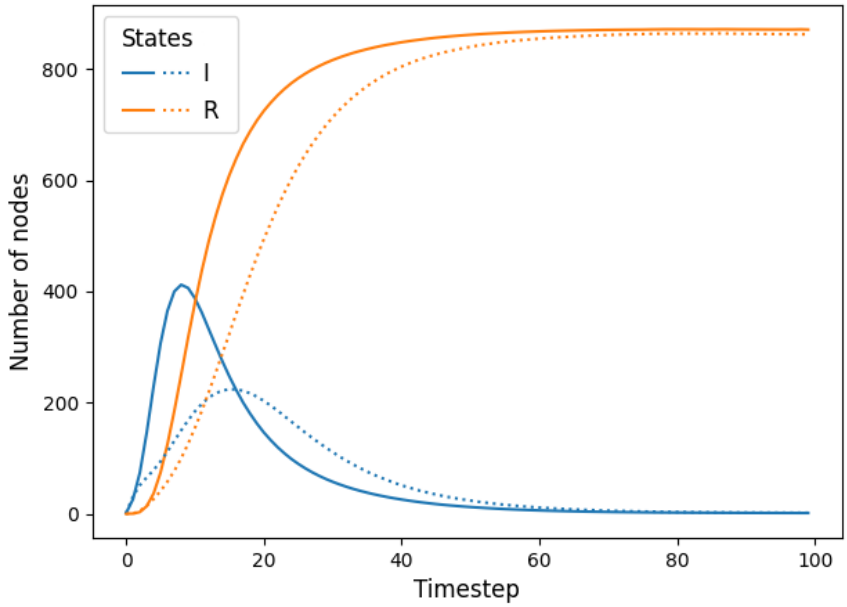

Figure 5.5: Progression over time of the number of nodes per state (spreaders and stiflers) with random vaccination and hub misinformation, where both start with a single node. The solid lines represents BA states and the dotted lines represent LFR. The data was obtained with infection rate $\beta = 0.5$ and vaccination rate $\theta = 0.2$, stifling rate $\gamma = 0.1$, forgetting rate $\delta = 0.1$, vaccination stifling rate $\zeta = 0$, averaged over 10,000 repeats.

in Figure 5.6. Overall we get lower TM values, as expected given the advantage to vaccination. BA and LFR continue with similar results between themselves, while DMS shows much lower TM overall, in comparison. On Figure 5.7, we can see a close up of timescale 4, and see that the difference between the networks still remains.

DMS

Again we see the effects of choosing a core node (or near the core) for spreading information. Since it cannot become stifled, vaccination can spread indefinitely until misinformation dies out or there are no more convertable nodes in reach. With this, starting in the center of the hierarchy allows it to spread much more effectively than misinformation. With the addition of a vaccinated stifler, it is likely we would see similar (reversed) values to when misinformation started in a hub. Despite this, with low timescales and low vaccination rate θ it will not spread very quickly. With a higher θ, the effects of a lower timescale can be reduced to a degree, but as infection rate β increases, its ability to contain misinformation decreases, but still within reasonable levels. As the timescale increases, it grows increasingly faster and will inevitably block all of the paths that misinformation could inevitably reach.

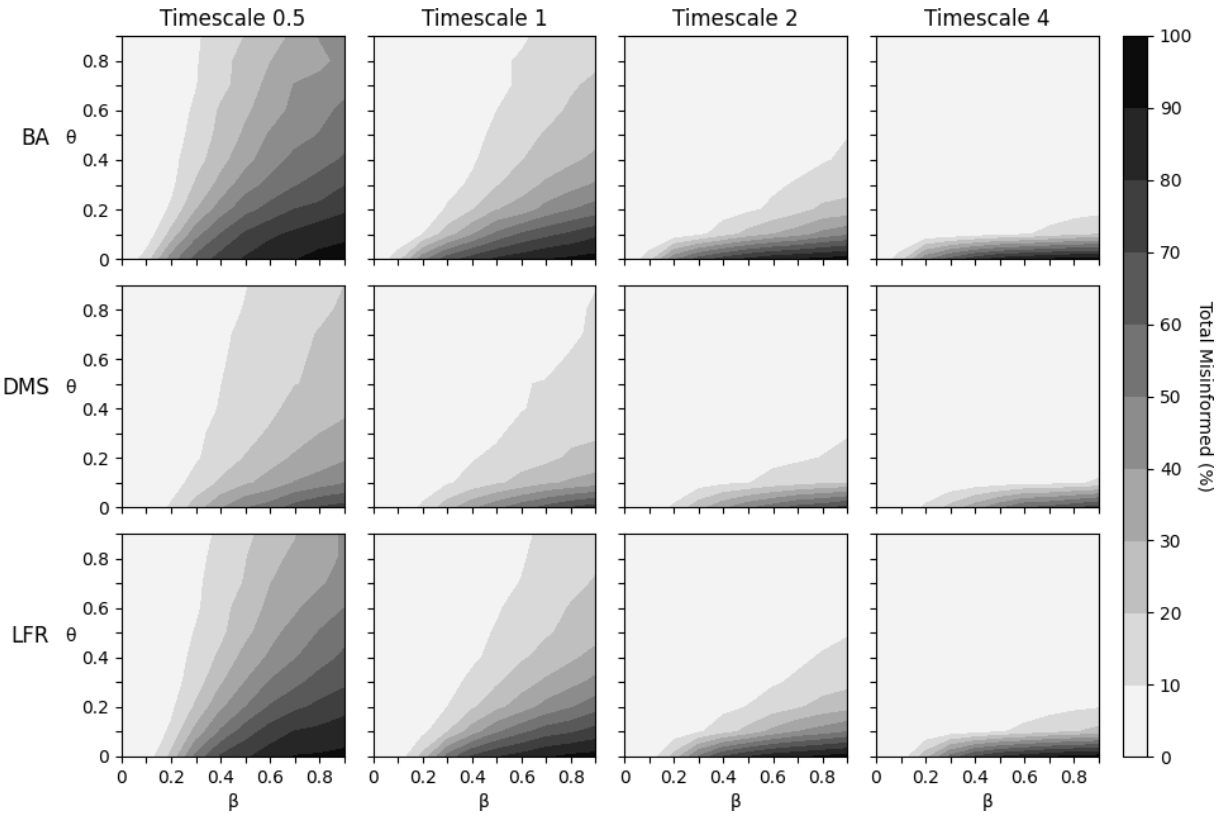

Figure 5.6: Contour graphs for Total Misinformed with different networks (rows: BA, DMS, and LFR) and timescales (columns: 0.5, 1, 2, and 4). Both the infection rate β and vaccination rate θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1. Misinformation starts in a random node while vaccination starts in the largest hub.

BA and LFR

Given the indefinite spread of vaccination and that both communities facilitate the spread of information, as seen before, by starting it in the largest hub, it reduces the spread of misinformation to a much higher degree. Except on higher β, but by increasing the timescale we can nullify its effect. Following what we saw when misinformation started in a hub with the growth of spreaders being the same between BA and LFR for a few time steps, we can expect the same to happen here with the growth of vaccinated individuals (see Figure 5.8). It should be noted that, for timescale 1 here, we get similar results to timescale 2 when both vaccination/misinformation are random (see Figure 5.2).

5.2.4 Polarisation

While the Total Misinformed can show us how effective increasing timescale or β is, it does not paint the full picture of what happens in the networks. One thing we can look at, to help our insight into how the spread of both informations characterise the networks, is polarisation. By this we mean to measure how polarised (or segregated) individuals and communities are if they have a tendency to be neighbours with individuals who share the same opinion (information in our case). The higher that tendency is, the more polarised a network is, the more segregated individuals are. In order to study this, we devised a

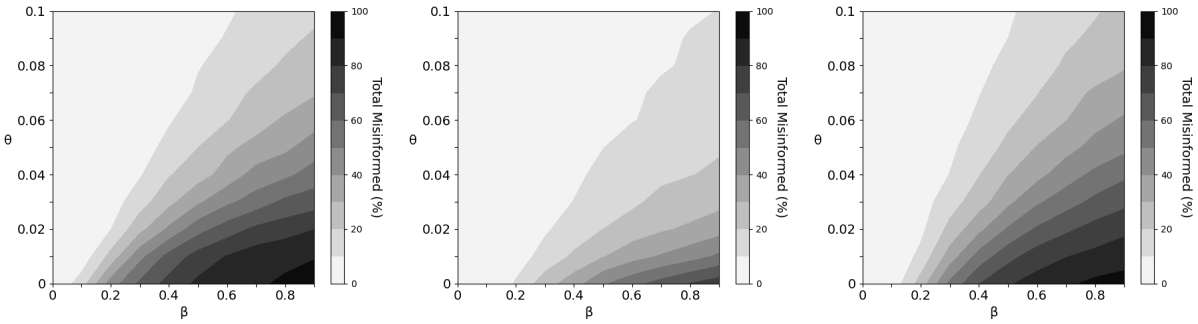

(a) (b) (c)

Figure 5.7: Contour graphs depicting how Hub vaccination and Random misinformation affects the various networks (a) BA, (b) DMS, (c) LFR, with a closer look at the interval 0 to 0.1 for vaccination rate θ, with timescale 4.

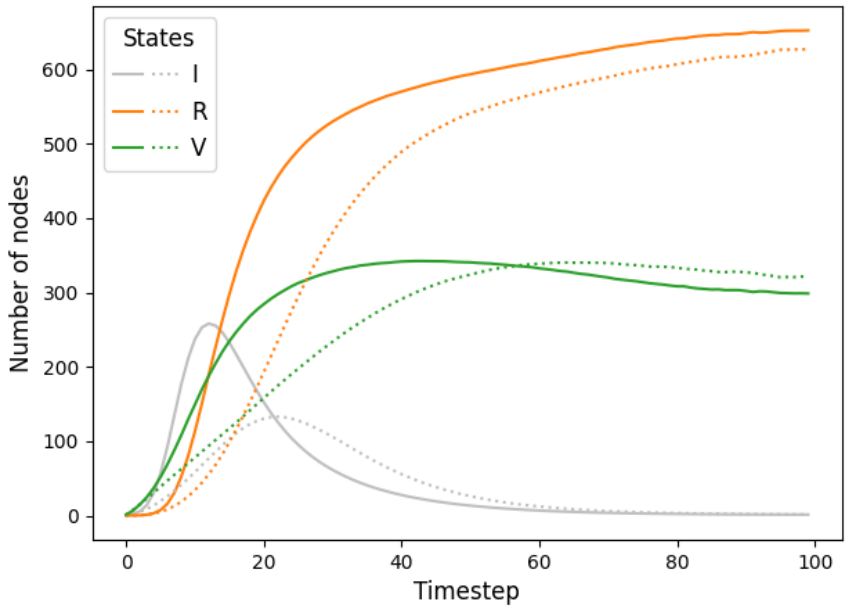

Figure 5.8: Progression over time of the number of nodes per state (spreaders, stiflers, and vaccinated) with hub vaccination and random misinformation, where both start with a single node. The solid lines represents BA states and the dotted lines represent LFR. The data was obtained with infection rate $\beta = 0.5$ and vaccination rate $\theta = 0.2$, stifling rate $\gamma = 0.1$, forgetting rate $\delta = 0.1$, vaccination stifling rate $\zeta = 0$, averaged over 10,000 repeats.

simple polarisation index P.

$$P = 1 - mean(\frac{\text{cross edges}}{\text{total edges}}, \text{per community}) \tag{5.1}$$

where cross edges are edges between nodes of different opinions (considering ignorants, stiflers, and vaccinated). This index is focused on communities, but since BA and DMS don't have communities per

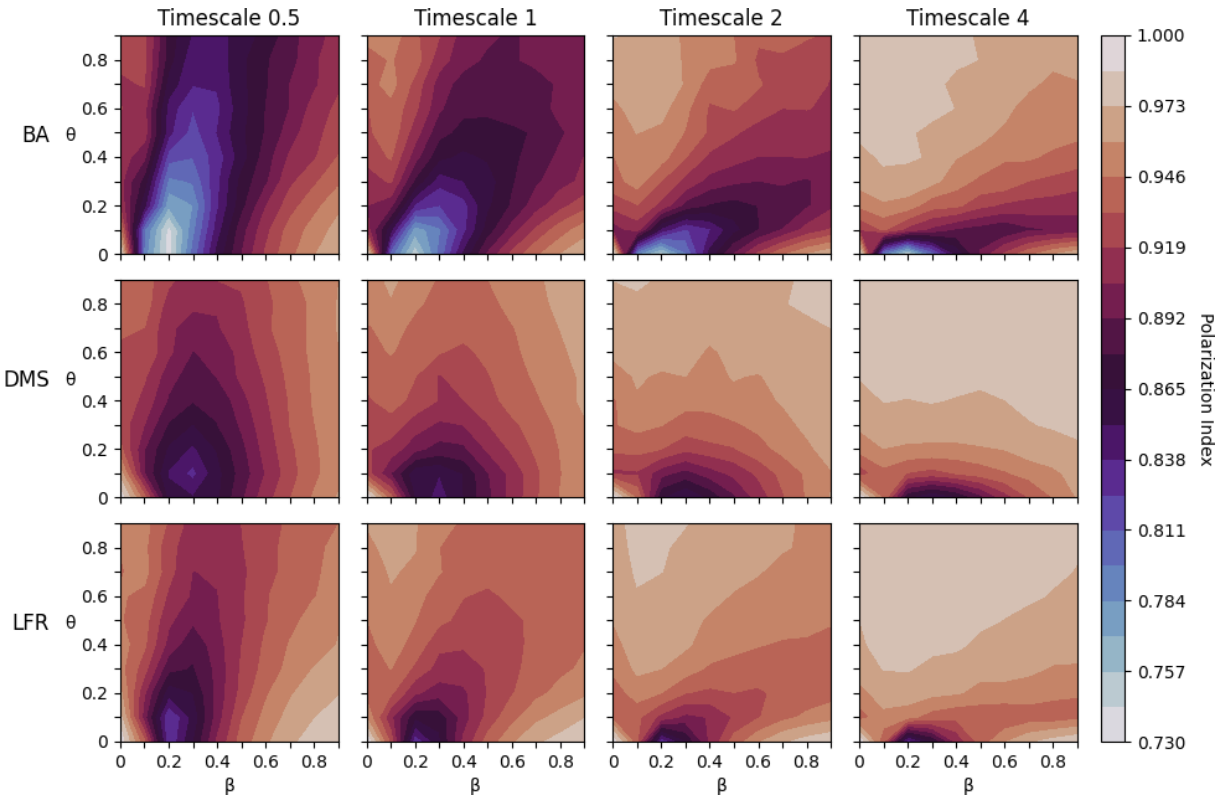

Figure 5.9: Contour graphs for the polarisation index with various networks (rows: BA, LFR, and DMS) and timescale values (columns: 0.5, 1, 2, and 4). Both infection rate β and vaccination rate θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1. Both misinformation and vaccination start in random nodes.

se, we consider them as a single community, where as in LFR we use the designated communities by the generation algorithm. What this index tells us is the polarisation average per community. The closer the value is to 1, the more homogeneous are the communities (or community).

Random Vaccination and Random Misinformation

With this index applied at the end of each simulation, we can rerun the simulation batches of Figure 5.2, obtaining the results in Figure 5.9. We can see that BA generally has a lower index, given that we considered the network a single community. The three possible states (neutral, informed and misinformed) will have more interactions between them due to the clusterless structure that BA has, leading to a more heterogeneous community, as seen by an example network in Figure 5.10a. On the other hand, LFR and DMS show much less polarisation overall. On LFR, the community structure helps to keep the three opinions mostly separated inside their communities, as seen in Figure 5.10b. Since the spread between communities takes longer and is more difficult, both spreaders and vaccinated tend to take over the communities that they reach first. And it is also possible, although dependent on the parameters, to have entire communities of susceptibles that remain untouched. This can happen if the paths to these communities become stifled and neither vaccination or misinformation can reach them. For DMS, the hierarchical structure forces both misinformation and vaccination to meet only at the edges of the areas

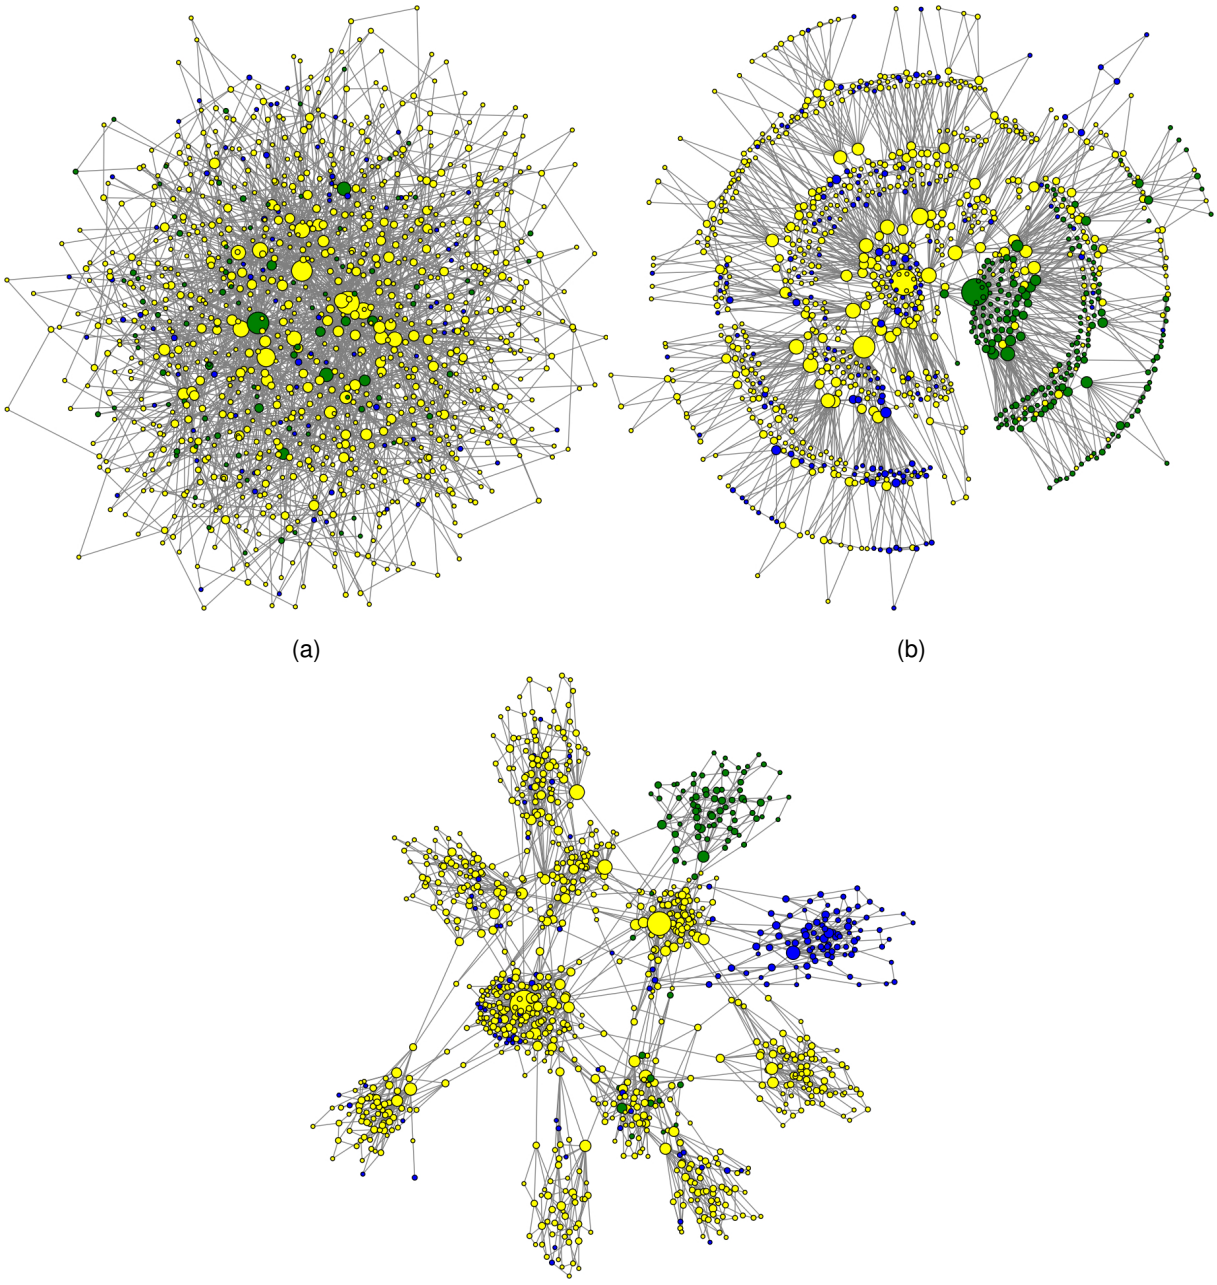

Figure 5.10: Examples of networks at the end of the simulation: (a) BA, (b) DMS, and (c) LFR. The colored nodes are: blue for ignorants, yellow for stiflers and green for vaccinated.

that they control, leading to a lower number of cross edges between them. The susceptible that remain on the network will depend on the parameters. Given this, we can see a clear difference in the shape of change that BA and LFR have compared to DMS, and we will start by discussing the former.

BA and LFR. We'll be focusing first on the results for timescale 0.5. The results can clearly be divided into three zones regarding β values. For 0 and very low β values, the index takes higher values as communities will be mostly made up of ignorants and vaccinated nodes, since misinformation doesn't spread much due to low β.

As we increase β, the networks enter into a more chaotic zone where misinformation becomes stronger in spreading but still not strong enough to overcome vaccination or to convert most susceptibles. This leads to more fights for control between misinformation and vaccination, while ignorants remain a significant part of the networks, leading to higher heterogeneity in the communities. In this zone, for low θ, vaccination is much more sporadic in its spreading, as is the same for misinformation, so that each side doesn't grow consistently on all fronts but only on occasional places where it manages to spread. This increases the area of cross contact between different opinions much more, when compared to higher θ or β. This area of cross contact is even larger in BA, because with low clustering, when it manages to spread, it will come in contact with mostly new nodes that it wasn't in contact with before. While in LFR, since nodes are mostly connected to other nodes in the same community (higher clustering), the area of cross contact will be be smaller since when it is able to spread, most likely the new nodes will be in the same community and possibly already were in contact before. As θ increases, the vaccination will spread much more consistently across its area of cross contact, and will lower the amount of cross edges. But still to a lesser degree in BA when compared to LFR. Vaccination will also reach more nodes much faster, but since the simulation stops when spreaders becomes 0, we still get lower polarisation values. Were it left to keep vaccinating until it couldn't anymore, the index would, we assume, be higher in this area.

As we increase β further into higher values, we enter another phase of high polarisation. The β allows misinformation to reach and convert communities much faster and efficiently, leading to less ignorants and cross-edges. Contrary to the chaotic phase, as we increase θ, polarisation goes down. For low θ misinformation can reach most of the nodes in the networks, while with higher θ, both sides are competing more to reach every node. And, in the case of LFR, both will convert entire communities (or almost fully) before the opposing information can reach them. The number of susceptible will also decrease as both sides have the ability to convert most of the ignorants they come in contact with. In BA with a single community, while both sides will grow fast and effectively, their area of contact will still be high because of the topology.

If we increase the timescale, there is a growth in the initial phase that causes a shift to the right of the chaos phase. This is due to the increased effect of θ with higher timescales. The zones with higher θ will become more and more polarised since most of the network/communities will end up being dominated by vaccination.

Regarding the values for infection rate $\beta = 0$, on the higher timescales, there is a very noticeable "harsh" drop to the next data point at $\beta = 0.1$. This is because on $\beta = 0$, the single misinformation node will not spread and the simulation will only stop when it transitions to a stifler via the forgetting rate δ, which is fixed for 0.1. Since this can take some steps until it the transition is successful, during this time the vaccination can spread freely leading to lower polarisation values. When β becomes 0.1, the misinformation will die out faster because of stifling and the vaccination will spread less, leading to more polarisation.

It should be noted that on timescale equal to 1, the central line of the polarisation originating from the lowest valued area gives roughly an identity line. We assume that this means that both informations

have the same spreading strength as their respective variable increases, when they spread at the same speed.

DMS. DMS gives us a different shape of the results, but the overall values are nearer to LFR than to BA. The minimum valued area in LFR and DMS is smaller in the former, but the rest of the lower valued areas are larger in the latter. The shift to the right that happens when timescale is increased is also no longer present, but instead the lower valued zones shrink down vertically. We assume that the main reason that the shape is different is because of not having the higher polarisation zone for low vaccination rate θ and high β that we see for BA/LFR. Focusing on timescale 0.5, for the 0 and low β phase, we get a similar behaviour with the vaccinated and ignorants making up most of the network, leading to a higher index.

As we move into the chaos phase, misinformation spreads more, and networks become more heterogeneous. But we still get lower values than LFR overall. This is due to the core and branch like structure, because the area of cross contact between the three sides will be higher than LFR. Since β is not high, misinformation doesn't spread consistently over the hierarchy and triangles, which increases cross contact.

Increasing the β further we reach a behaviour different to BA/LFR. The tendency is not longer for the polarisation index to increase into the $\theta = 0$ and $\beta = 0.9$ zone, at least not as strong. This tendency disappears due to the DMS topology. Regardless of whether vaccination spreads slow, fast or not at all, the contact area between both sides will mostly be the same, because of the triangle hierarchy. They will only have cross edges at the beginnings of branches or sub-branches, which have a very limited area of contact. The rest of the network will be cutoff by those areas of contact and left for the information that controls each side of that area to spread freely. Since stifling can also prevent misinformation from accessing a certain area, this will also increase the polarisation, albeit to a smaller degree.

Overall we see a general tendency for polarisation to increase the further we get from the lowest point. As timescale increases, we see a downwards shift instead as for higher θ the networks will be mostly made up of vaccinated individuals, because they will spread effectively and consistently, reducing the cross contact area even further. We assume the shift to the right no longer happens due to the topology effect on the higher β discussed above. We also see the same effect between $\beta = 0$ and $\beta = 0.1$ that we saw in BA and LFR.

Hub Misinformation

Starting misinformation in the largest hub we get the polarisation results seen in Figure 5.11. The zones remain similar, but the effect of θ is reduced giving the impression that every area is stretched upwards. This reduction is due to the advantage that misinformation has by starting in a hub. The lower limit of polarisation decreases, in both BA and DMS, while in LFR the limits stay roughly the same. This means that the former can reach a much more heterogeneous state, depending on the variables.

BA and LFR. Focusing on timescale 0.5, the very low β zone stays roughly the same, both in reasoning and results, since even though misinformation starts in a hub it will spread very little or nothing. The chaos zone is much more pronounced upwards, since even with high θ, where vaccination spreads

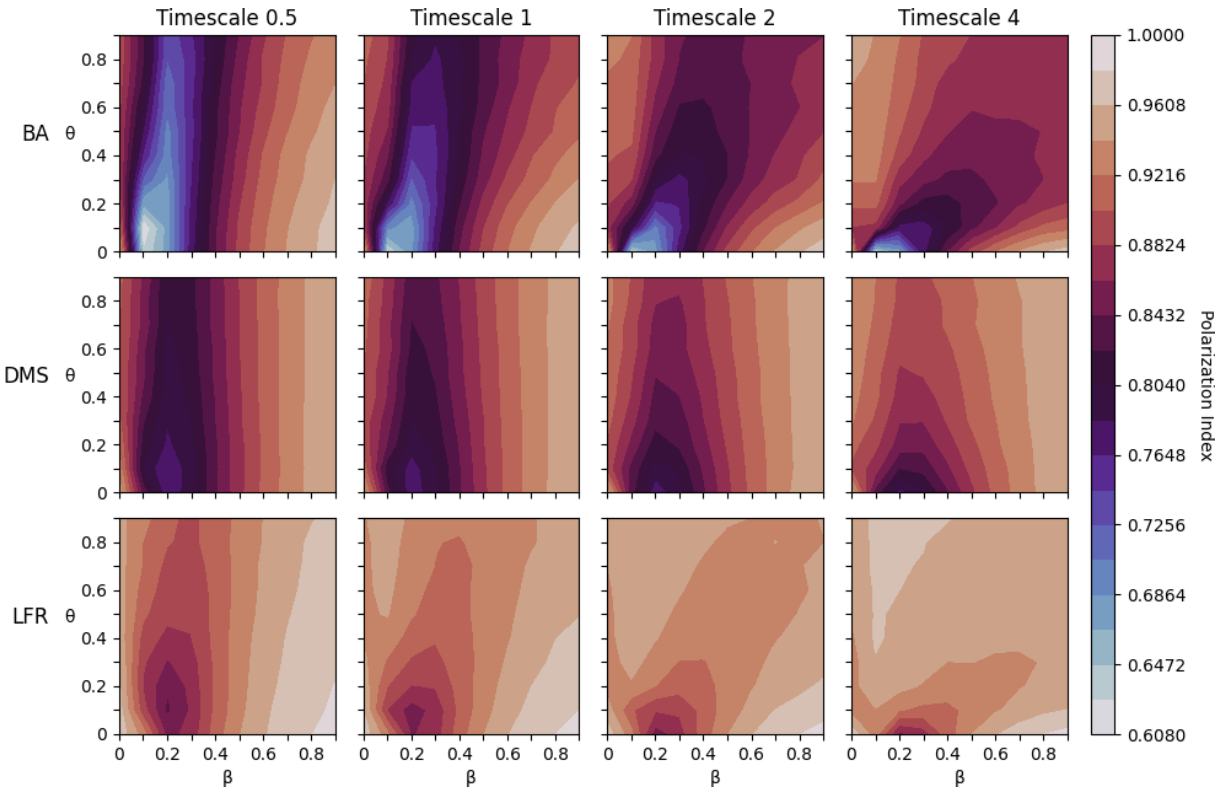

Figure 5.11: Contour graphs for the polarisation index with various networks (rows: BA, LFR, and DMS) and timescale values (columns: 0.5, 1, 2, and 4). Both the infection rate β and vaccination rate θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1. Vaccination starts randomly while misinformation starts in the largest hub.

a lot, misinformation will spread more than before (starting in random). This leads to a decrease in the rise of polarisation as θ increases, causing the noticeable upwards change. The center of the chaos zone (i.e. β value for the center of the lowest valued area) for BA change approximately from 0.2 to 0.1 β. The same does not happen for LFR, which maintains the same center. We assume this is related to the reason why polarisation limits do not change in LFR, but do in BA. Although this will be not addressed here. On the high β zone, we also see this change, since starting in a hub increases considerably the effect of β. Comparing it to the results in Figure 5.4 for timescale 0.5, θ has a very little effect on TM in the chaos zone, similarly to here. Furthermore, on the high β zone, misinformation will rule most of the network, which matches with the high polarisation here. As we increase timescale, like before, there is a shift to the right of the chaos zone as θ's effect is boosted and the low β zone expands. Nonetheless, even with high timescales, misinformation is still able to resist misinformation considerably, and the overall polarisation does not increase as much as when both start randomly.

DMS. Again focusing on timescale 0.5, all three zones extend upwards, with the very low and high β zones remaining roughly the same. There is a shift in the chaos zone center of approximately 0.3 to 0.2 β, similar to BA as discussed above. Other than that and the upward extension due to starting in a hub, the chaos zone behaves similarly, except for when we increase timescales. Likewise to BA and LFR, the shift (in this case downwards) becomes softer due to the increased strength of misinformation.

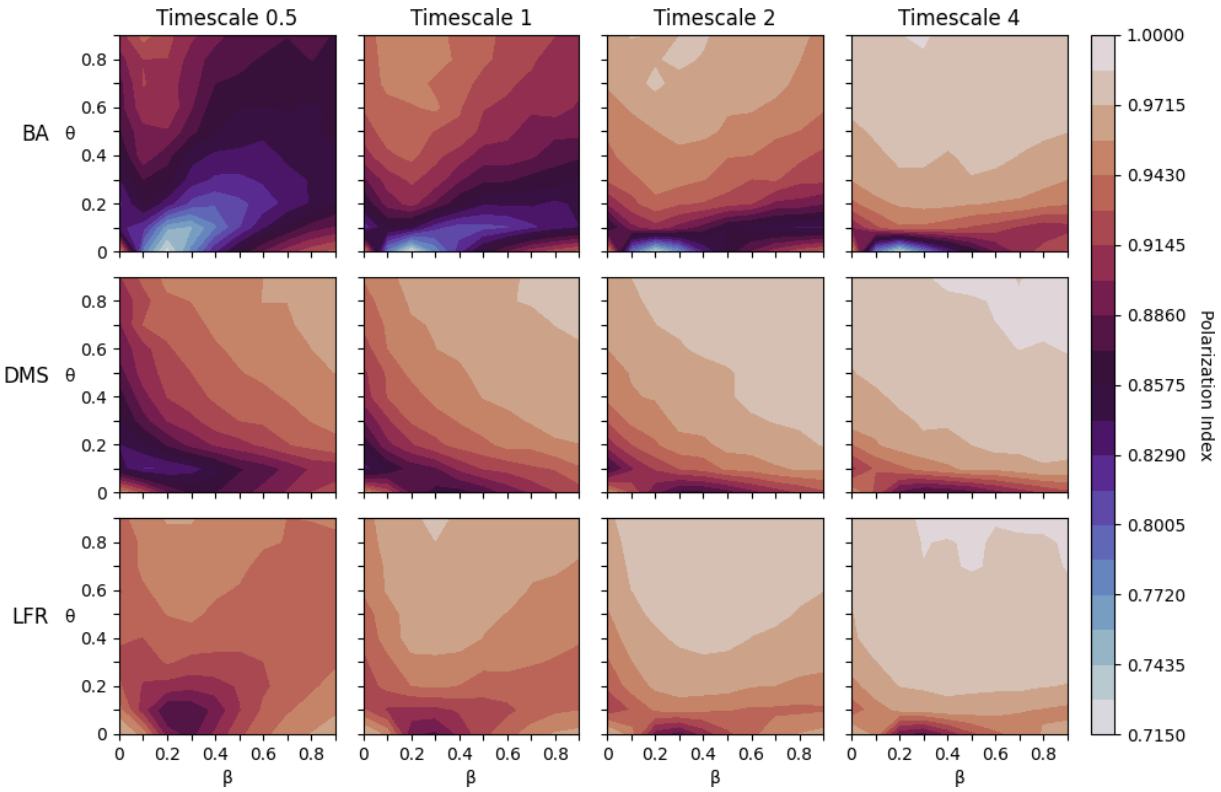

Figure 5.12: Contour graphs for the polarisation index with various networks (rows: BA, LFR, and DMS) and timescale values (columns: 0.5, 1, 2, and 4). Both the infection rate β and vaccination rate θ vary from 0 to 0.9 with step 0.1. Vaccination stifling rate ζ is set to 0, while stifling rate γ and forgetting rate δ are set to 0.1. Misinformation starts randomly while vaccination starts in the largest hub.

Hub Vaccination

With vaccination starting in the largest hub, we get the results seen in Figure 5.12. With the exception of the very low β zone, the rest of the zones for BA and LFR remain somewhat the same, albeit with a larger shift to right. However, DMS gives us completely different zones. With this type of vaccination, the overall interval for polarisation remains similar to when both informations started randomly. The center for BA remains the same while LFR's increases slightly. For DMS the center isn't as clear due to the horizontal area of the lowest value.

BA and LFR. Focusing on timescale 0.5, the reason why the very low β zone is much less polarised than before is due to misinformation not spreading or spreading very little. With $\beta = 0$, until the single spreader forgets, the vaccination starting in a hub is able to spread more, leading to less polarisation between vaccinated and ignorants. As vaccination rate θ increases, the spread becomes more consistent and the tendency is for polarisation to increase. As β increases, this tendency happens much faster since vaccination has more time to spread. This continues into the chaos zone which is now shifted considerably due to the increased vaccination strength. As β increases into the higher values, the upper echelons of θ will become less polarised as misinformation will be able to compete with vaccination. As timescales increase, we see the same shift due to the enhanced effect of increasing θ, but much more pronounced due to the hub start.

DMS. Focusing on timescale 0.5, the same logic for the very low β zone applies. But for the two other zones, it becomes more difficult to explain via the infection rate β. We can instead look at the graph as composed of two θ zones. On low θ (chaos), due to starting in the core (or near it) it will spread enough to have a considerable area of contact with ignorants and spreaders, since the spread is not consistent, leading to lower polarisation. As β increases, misinformation will convert more and more of the branch where it started, given that it's cutoff from the core by vaccination. This will lead to less areas of contact between spreaders and the rest, leading to less polarisation. Furthermore, since misinformation will take longer to die out, vaccination can spread more in the other branches, therefore reducing the ignorant population further and increasing polarisation. As θ increases, it's able to spread further and more consistently from the core, therefore reducing the population of ignorants and the space for misinformation to spread (until it is blocked by it), and increasing polarisation considerably. As β grows, misinformation will spread more in its branch, and the amount of cross edges will be reduced greatly or even completely. As timescale increases, vaccination will spread even faster from the core, out performing misinformation. It is possible that it may even envelope the spreader node before it has the chance to block the branch for itself.

It should be noted, that both extremes (hub misinformation and hub vaccination) complement each other. This is easy to see on timescale 0.5 for BA and LFR, as the vertical chaos zone and the more polarised zones in hub misinformation are roughly switched when we look at hub vaccination.

5.2.5 Concluding Remarks

Timescales give us the opportunity to study how two informations with similar or different speeds spread and affect each other.

We discovered that BA and LFR (our subset of possible LFRs) have extremely similar results for Total Misinformed (TM), independent of where either information starts. Despite leading to similar TM at the end of our simulations, where LFR always has the lower value, LFR has a different spreading pattern due to its restrictive structure. This raises the implication that both networks share some implicit similarity in their structure despite their differences. The initial similarity in the spreaders quantity when misinformation starts in a hub could provide some insight into it. Furthermore, the structure of each community and the topology of communities and how they connect with each other might provide further insight. Despite this similarity, LFR networks lead to much more polarised networks due to the community structure which helps to maintain the information "trapped" inside the community and keep out new information. BA's chaotic structure enforces more heterogeneity since information can freely propagate.

In both networks, vaccination is very weak when it's slower than misinformation (see timescale 0.5), especially against virulent rumours, as would be expected. As we match the speed (in case of hub vaccination) or double it (in case of both random) or quadruple it (in case of hub misinformation) we can see great results in reducing misinformation spread. Given the situation, in order to combat it effectively we need to know where it started and its spreading strength (similarly to R_0 in epidemiology).

With DMS, we saw that hierarchical networks can generally resist the spread of misinformation but also the spread of vaccination, so it's a double edged sword. However, when either start at the top of the

hierarchy, their spreading potential is amplified tremendously, regardless of other variables. In the case of misinformation, even with very high timescales, its extremely difficult to reduce misinformation spread significantly, due to its restrictive structure. Otherwise, increasing the timescale provides great reduction of misinformation, but past a certain point the structure forces diminishing returns.

When we look at polarisation across all three starting scenarios, the same order of overall polarisation is maintained. BA is the least polarised (more heterogeneity) and LFR is the most. LFR's communities cause information bubbles due to the difficulty and time involved in spreading beyond a community. BA on the other hand, with it's single community with null clustering, leads to both informations spreading wildly without staying near to where they were before, which leads to a high area of cross contact. DMS's hierarchical structure also causes information bubbles due to the easily blocked paths for propagation, but the polarisation is lower than LFR, because the structure leads to a higher area of cross contact. In LFR this area is minimised because of the clustered communities and low inter-community edges.

From these differences in polarisation, we can make a generalised assumption: increase in order (i.e. increase in average clustering coefficient) does not correlate with an increase in polarisation. A better metric for a simple correlation with polarisation might be Average Path Length. Remembering the metrics from Table 3.2, BA had on average $\sim$4 steps between every node, DMS had $\sim$4.9 and LFR $\sim$6.8. These match the order we see for polarisation increase, and also the amount of difference between them (except with both random starting positions where DMS and LFR are close, but LFR has more polarisation overall). As nodes become more distant, on average, from each other, the longer information will take to reach the various nodes in the network. By taking more time, information will spread more in the vicinity of where it started and less in the faraway nodes of the network. This will create clusters of the same information, leading to more polarisation. More distance also means more time to spread consistently where it already is (nearer to the start). In LFR with a high enough infection rate β or vaccination rate θ, the starting community will likely be almost fully infected before the information even reaches the furthest community.

With regards to how polarisation changes with variations to β and θ, both BA and LFR gave similar change patterns despite the differences in overall polarisation and the network structures. This continues the trend of similarity seen in the TM results. DMS gives us a different polarisation pattern, because of the effects of structure on the propagation and how nodes are connected which affect the area of cross contact.

There is no generalised conclusion on whether increasing θ or β, separately, leads to increase or decrease in polarisation, since it depends on the network, the value of the other variable and the timescale.

5.3 Vaccination Stifling

With a spreading vaccination we can now get results on how combining it with stifling influences the misinformation spread.

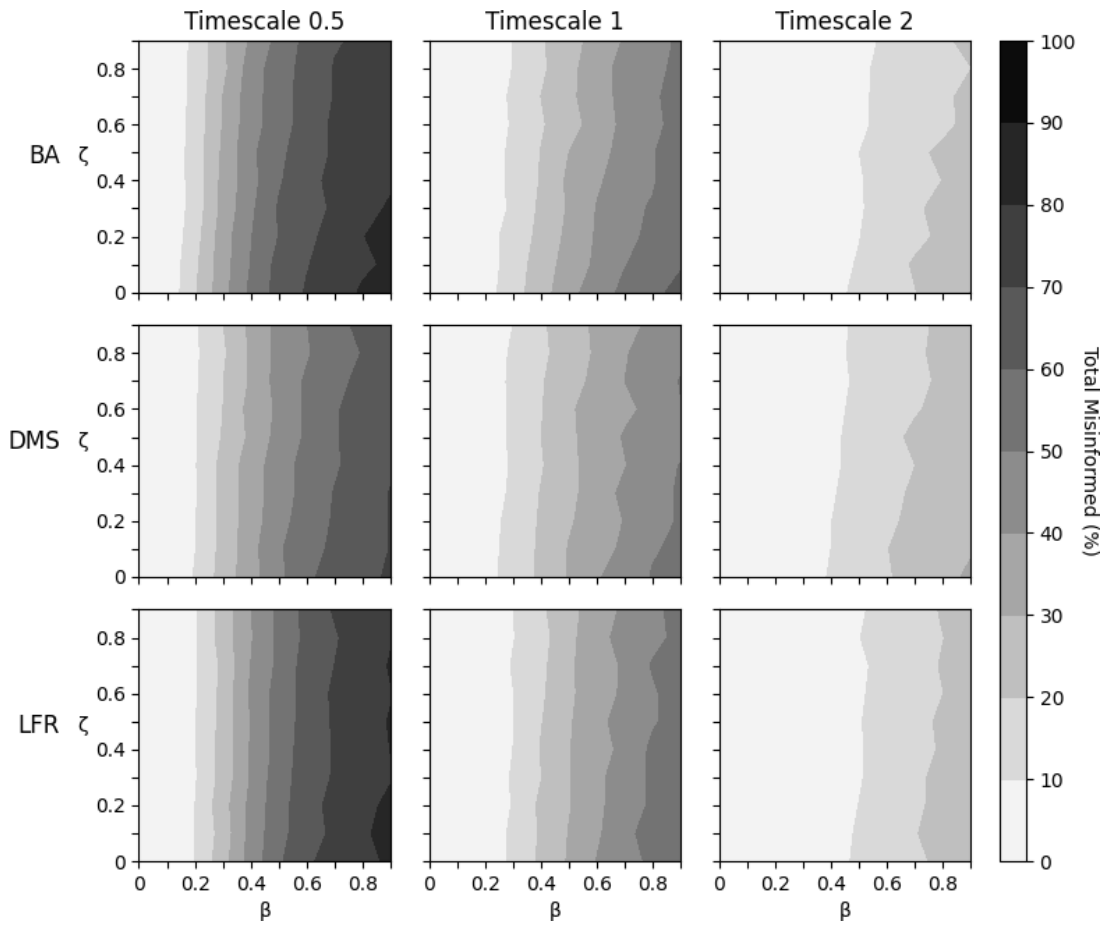

Figure 5.13: Contour graphs for Total Misinformed with vaccination stifling, with different networks (rows: BA, DMS, and LFR) and timescales (columns: 0.5, 1, 2, and 4). Both the infection rate β and stifling vaccination rate ζ vary from 0 to 0.9 with step 0.1, while the vaccination rate θ is fixed at 0.5, and stifling rate γ and forgetting rate δ are fixed at 0.1. Both vaccination and misinformation start in a single random node.

5.3.1 Random Vaccination and Misinformation

Starting with both rumours beginning in a random node, we fixed vaccination rate θ at 0.5 and instead varied the vaccination stifling rate ζ, and obtained the results seen in Figure 5.13.

Despite DMS having lower TM overall, the behaviour seems consistent for the three networks, especially on lower infection rates (β), where increasing ζ gives negligible returns. Even as timescale increases, this is maintained. It is only on higher β that we see any actual change in TM as ζ increases. Owing to higher β, misinformation and vaccination spread more and have increased contact with each other, which will lead to more stifling interactions and a higher chance for them to affect TM. The more successful these events are, via ζ, the more TM is affected. With lower β and/or θ, both sides will interact less, making stifling irrelevant.

Both sides meeting during their infancy is important. Stifling one misinformed when misinformation makes up a low percentage of the network is more important than if it made up a high percentage. This is because if that one node is not stifled, then it spreads and creates a chain reaction of spreading

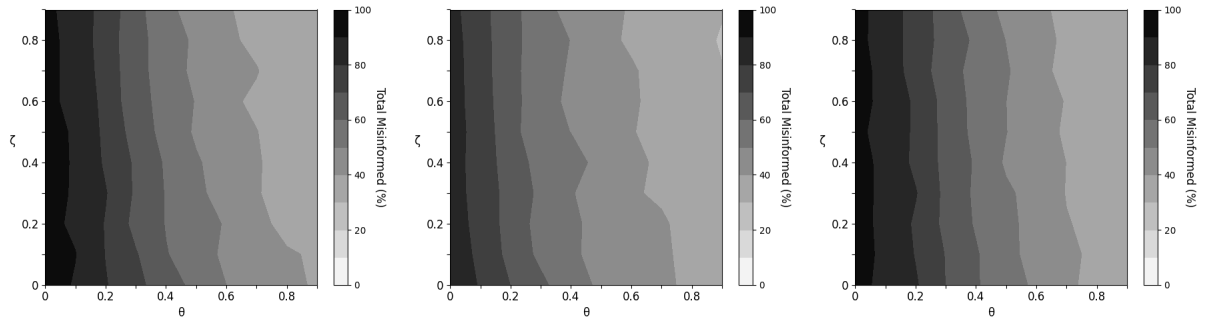

(a) (b) (c)

Figure 5.14: Contour graphs for Total Misinformed with vaccination stifling, with different networks (a) BA, (b) DMS, and (c) LFR. Both the vaccination rate θ and stifling vaccination rate ζ vary from 0 to 0.8 with step 0.1, while the infection rate β is fixed for 0.8, stifling rate γ and forgetting rate δ are fixed at 0.1, and timescale is 1,

that could lead to a high number of spreaders. Of course this depends on where that node is in the network, if it has neighbours, and how the topology of the network is. On DMS, unless the stifled node is connected to branches with a lot of untouched nodes, stifling it or not won't matter much, since it will be mostly blocked between the zone of the misinformation where it became a spreader from and the zone of vaccinated that might stifle it. On BA, excluding lower degree nodes, since they will also be mostly blocked on both sides, if the misinformed are not stifled, they can spread to a new set of nodes outside of its neighbourhood and lead to a possible cascading effect of new spreaders. With BA topology both sides will also meet on more fronts, increasing the effect of stifling as seen from the results.

On LFR, we see a behaviour different to BA. Even by increasing β and/or timescale, stifling shows little to no effect. We assume that this is due to the topology, and the retardation of the spread of both sides (from the topology and starting in a random node), leading to them meeting later (on average) and significantly lowering the effect stifling has. Also, since communities become very polarised in LFR, both sides will also interact less, leading to another reduction in its effect.

With a higher fixed θ, we might see larger changes given ζ as both sides will meet sooner. But as vaccination grows larger, faster than misinformation, the interactions between both sides will decrease and the stifling effect will lower again. Given this, it would be expected that the gain to stifling effect from increasing ζ would become lower as β reaches higher levels. At the same timescale, what we see when we fix β at 0.8 is that the effect appears to keep increasing, with BA displaying the largest increases, as seen in Figure 5.14. The lines between the lower valued zones show very similar slopes in their boundary lines, so the increase might have plateaued.

5.3.2 Hub Vaccination

If we instead start vaccination in a hub, the effect of stifling increases and becomes more noticeable even for lower β, as seen in Figure 5.15. Vaccination spreads more and faster, causing stifling interactions to happen more and sooner, leading to a larger reductions in TM. While BA and DMS have an

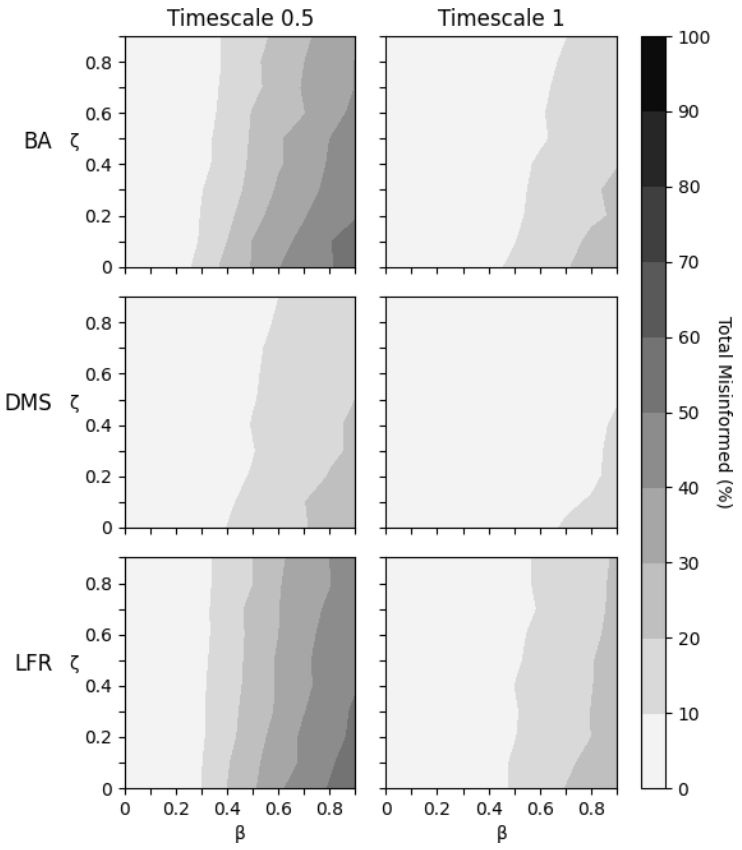

Figure 5.15: Contour graphs for Total Misinformed with vaccination stifling, with different networks (rows: BA, DMS, and LFR) and timescales (columns: 0.5, 1, 2, and 4). Both the infection rate β and stifling vaccination rate ζ vary from 0 to 0.9 with step 0.1, while the vaccination rate θ is fixed at 0.5, and stifling rate γ and forgetting rate δ are fixed at 0.1. Vaccination starts in the largest hub while misinformation in a random node.

increased effect overall, in LFR it remains the same for lower β. Yet, for higher β the effect increases slightly. We assume that vaccination starting in a hub, leaving its initial community faster and therefore reaching misinformation faster, will generate more stifling interactions and lower the spread of misinformation. Since θ is fixed at 0.5, the spread of vaccination is also not generally consistent so the gaps in its spread, that could be used by misinformation, will not be used as much for it to spread due to stifling. Overall, it seems stifling can help to fight misinformation independent of the network when we are dealing with a strong spreading vaccination.

5.3.3 Hub Misinformation

If misinformation starts in a hub, we see a larger effect for stifling the higher β is, when compared to hub vaccination or both starting in a random node, as seen in Figure 5.16. With misinformation starting in a hub, the amount of spread that it can reach is high, therefore the earlier and the more stifling that happens, the more TM will be affected. The higher the expected TM is with no stifling, the greater reduction in TM we will see as stifling increases. BA and DMS show larger changes at higher β as we increase ζ, whereas LFR still produces little change from stifling. These changes are exacerbated as

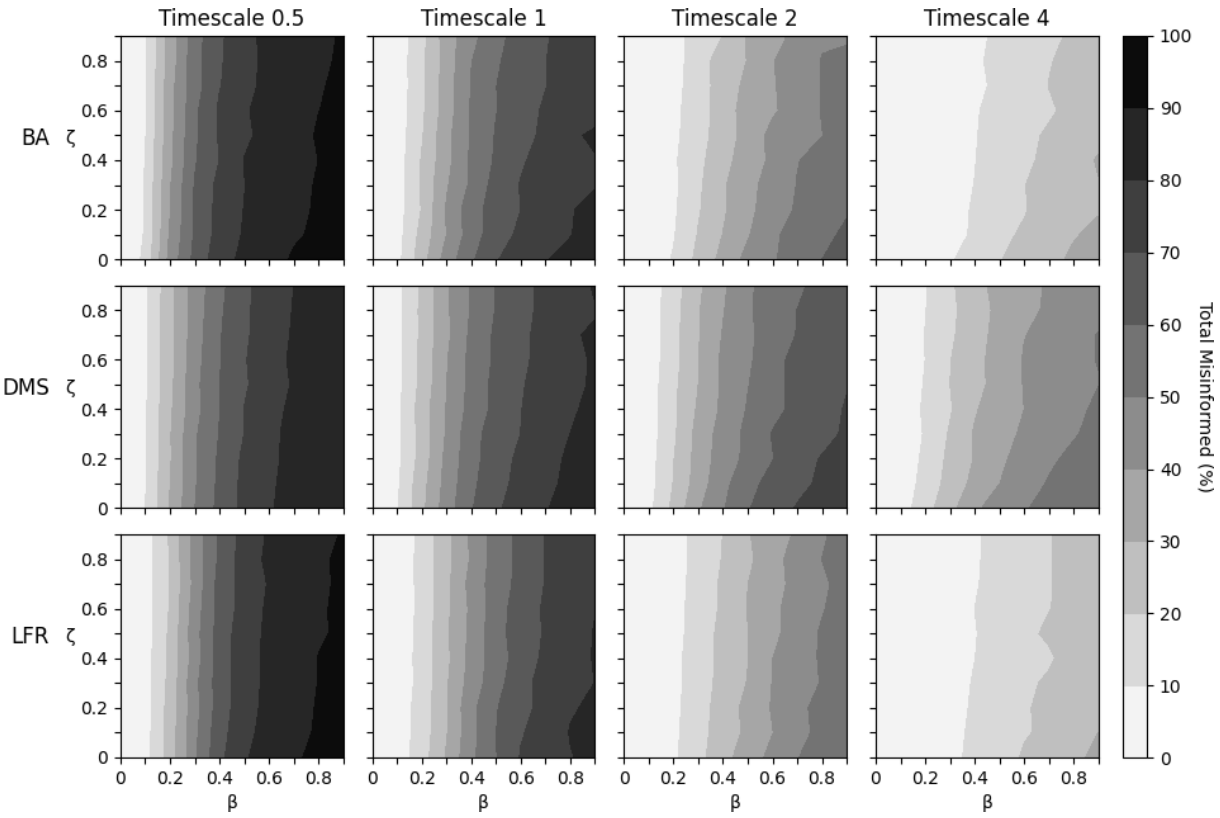

Figure 5.16: Contour graphs for Total Misinformed with vaccination stifling, with different networks (rows: BA, DMS, and LFR) and timescales (columns: 0.5, 1, 2, and 4). Both the infection rate β and stifling vaccination rate ζ vary from 0 to 0.9 with step 0.1, while the vaccination rate θ is fixed at 0.5, and stifling rate γ and forgetting rate δ are fixed at 0.1. Misinformation starts in the largest hub while vaccination in a random node.

timescale increases. The community structure of LFR continues to impede stifling from reducing TM significantly, while on BA and DMS the higher areas of cross contact aid stifling.

5.3.4 Concluding Remarks

In general, it appears that stifling is not helpful for obtaining large decreases in misinformation spread. Unless we are dealing with an extremely virulent misinformation (via infection rate β) or vaccination (either via vaccination rate θ or timescale), the TM losses are minimal. Only in those severe cases of high virulence, or where vaccination starts with a hub, do we get tangible reductions. In the case of both a virulent misinformation and vaccination, stifling also appears to be helpful without needing further study.

6 | Conclusions

This research allowed us to get a better insight into how structure can influence the spread of misinformation when the truth is introduced to fight it. Regardless of whether we are dealing with top-down (classical) or bottom-up (dynamic) vaccination, structure plays the most important role.

Our results showed that hierarchical networks can impact negatively the spread of both sides (resistance to change), when dealing with a bottom-up vaccination. While, on a top-down approach, the structure provides us with a potent effect in lowering misinformation, since both Acquaintance and Hub strategies provide low values of misinformation with very small vaccination fractions. However, on community and low clustering structures, we don't see this similarity between strategies. Both of these structures give similar results, albeit marginally lower with communities, but with a large difference between Acquaintance and Hub, so Hub is better if we can get global information on the network. Otherwise we require a much larger vaccination fraction to get similar results. On dynamic vaccination, both of these structures behave similarly, with the exception of more fringe cases with vaccination stifling. That being so, community structures that are sparsely connected slow down the spread of misinformation, but ultimately give us the same results as low clustering networks. However, sparsely connected community structures lead to substantially more polarisation, meaning individuals inside the same community mostly share the same opinion.

We also saw how these structures behave under targeted attacks (misinformation starting in the largest hub) and targeted defences (vaccination in the largest hub). While targeted attacks produce overall, more misinformation in the networks, we can still combat it if the structure allows it. Whereas low clustering and community structures allow us to increase the timescale and vaccination efficacy to reduce misinformation spread, since hierarchical structures hinder the spread of information, it is increasingly difficult to combat misinformation as the misinformation becomes more viral, even with the increase of timescale and vaccination efficacy. On targeted defences though, these structures produce the lowest misinformation due to the same reasons, as misinformation cannot compete against the advantageous spread of the truth.

Whether or not clustering also plays a similarly important role is more difficult to answer, since separation between it and the structure is complicated. Given this, our results were inconclusive in this regard.

Another aspect that was studied was vaccinated individuals stifling misinformation. Based on our analysis it seems that stifling is only valuable when we are dealing with an overly viral misinformation. Vaccinated individuals that stifle spreaders are a costly endeavour (not only on a personal level but also on getting individuals to do it), so they are not very relevant in the case for most misinformation. Nevertheless, even an extremely viral misinformation returned negligible changes from any stifling, when tested with our community structures.

From this research, several questions emerged that could be the focus of future works. Firstly, with our three networks, we can test various types of structures but real social networks tend to be a mix of the qualities that each of these possess. They tend to have high clustering but are mostly not hierarchical, and individuals tend to form communities. Therefore, a real social network (or various examples) could be used with our model to get comparable results to the ones in our research, in order to see how our studied networks compare to a real one. Secondly, studying the triangle reduction of DMS networks which could provide an insight into the separation of structure and clustering. With the change to degree distribution and clustering, we expect that the misinformation results would be lower. How much lower they are comparably should show how important clustering is. Another step to studying this could be an intermediate mix, where the simulations take place in regular DMS networks, but any changes to states in a node would be applied to nodes in the triangle it belongs to (according to the reduced DMS network). Finally, given the similar results with BA and LFR, we can study if this behaviour maintains as we change the parameters of generating LFR networks. In particular, the fraction of inter-community edges, which changes how connected each community is to each other. As it increases, we believe that the community will become progressively more like a BA, since at fraction 1 no nodes will be connected with nodes within their community leading to very low or 0 clustering. Therefore, as it increases, it might keep behaving like a BA but with polarisation decreasing.

Improving the software solution is also a part of the future work to be done. Overall the system needs refining and optimising due to having been fast tracked to obtain results. The strategies system needs to be improved, as does the flow of information through the main system and into the modules. An event based simulation (as opposed to the implemented batch based system) should be added, and to allow the user to pick between the two. The configurability of the simulator, data processors, and other parts of the system need to be improved. An interesting metric to develop could be to know which transitions individuals used to reach their final states.

Bibliography

[1] Afassinou, K.: Analysis of the impact of education rate on the rumor spreading mechanism. Physica A: Statistical Mechanics and its Applications **414**, 43 – 52 (2014). https://doi.org/10.1016/j.physa.2014.07.041

[2] Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. Journal of economic perspectives **31**(2), 211–36 (2017)

[3] Barabási, A.L., Albert, R.: Emergence of scaling in random networks. science **286**(5439), 509–512 (1999)

[4] Bettencourt, L.M., Cintrón-Arias, A., Kaiser, D.I., Castillo-Chávez, C.: The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. Physica A: Statistical Mechanics and its Applications **364**, 513 – 536 (2006). https://doi.org/10.1016/j.physa.2005.08.083

[5] Cohen, R., Havlin, S., Ben-Avraham, D.: Efficient immunization strategies for computer networks and populations. Physical review letters **91**(24), 247901 (2003)

[6] Daley, D.J., Kendall, D.G.: Epidemics and rumours. Nature **204**(4963), 1118–1118 (1964)

[7] Dorogovtsev, S.N., Mendes, J.F., Samukhin, A.N.: Size-dependent degree distribution of a scale-free growing network. Physical Review E **63**(6), 062101 (2001)

[8] Ferrara, E., Varol, O., Davis, C., Menczer, F., Flammini, A.: The rise of social bots. Communications of the ACM **59**(7), 96–104 (2016)

[9] Fornito, A., Zalesky, A., Bullmore, E.: Fundamentals of brain network analysis. Academic Press (2016)

[10] Fruchterman, T.M., Reingold, E.M.: Graph drawing by force-directed placement. Software: Practice and experience **21**(11), 1129–1164 (1991)

[11] Gallos, L.K., Liljeros, F., Argyrakis, P., Bunde, A., Havlin, S.: Improving immunization strategies. Physical Review E **75**(4), 045104 (2007)

[12] Gilbert, E.N.: Random graphs. The Annals of Mathematical Statistics **30**(4), 1141–1144 (1959)

[13] Greenhalgh, D.: Optimal control of an epidemic by ring vaccination. Stochastic Models **2**(3), 339–363 (1986)

[14] Han, S., Zhuang, F., He, Q., Shi, Z., Ao, X.: Energy model for rumor propagation on social networks. Physica A: Statistical Mechanics and its Applications **394**, 99–109 (2014)

[15] Hartley, K., Vu, M.K.: Fighting fake news in the covid-19 era: policy insights from an equilibrium model. Policy Sciences **53**(4), 735–758 (2020)

[16] Huang, J., Jin, X.: Preventing rumor spreading on small-world networks. Journal of Systems Science and Complexity **24**(3), 449–456 (2011)

[17] Jin, F., Dougherty, E., Saraf, P., Cao, Y., Ramakrishnan, N.: Epidemiological modeling of news and rumors on twitter. In: Proceedings of the 7th workshop on social network mining and analysis. pp. 1–9 (2013)

[18] Kamada, T., Kawai, S., et al.: An algorithm for drawing general undirected graphs. Information processing letters **31**(1), 7–15 (1989)

[19] Kermack, W.O., McKendrick, A.G.: A contribution to the mathematical theory of epidemics. Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character **115**(772), 700–721 (1927)

[20] Kretzschmar, M., Van den Hof, S., Wallinga, J., Van Wijngaarden, J.: Ring vaccination and smallpox control. Emerging infectious diseases **10**(5), 832 (2004)

[21] Kucharski, A.J., Eggo, R.M., Watson, C.H., Camacho, A., Funk, S., Edmunds, W.J.: Effectiveness of ring vaccination as control strategy for ebola virus disease. Emerging infectious diseases **22**(1), 105 (2016)

[22] Lancichinetti, A., Fortunato, S., Radicchi, F.: Benchmark graphs for testing community detection algorithms. Physical review E **78**(4), 046110 (2008)

[23] Lazer, D.M., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., et al.: The science of fake news. Science **359**(6380), 1094–1096 (2018)

[24] Lelarge, M.: Efficient control of epidemics over random networks. ACM SIGMETRICS Performance Evaluation Review **37**(1), 1–12 (2009)

[25] Madar, N., Kalisky, T., Cohen, R., Ben-avraham, D., Havlin, S.: Immunization and epidemic dynamics in complex networks. The European Physical Journal B **38**(2), 269–276 (2004)

[26] Maki, D., Thompson, M.: Mathematical models and applications: with emphasis on the social, life, and management sciences. Prentice-Hall (1973)

[27] Martens, B., Aguiar, L., Gómez-Herrera, E., Mueller-Langer, F.: The digital transformation of news media and the rise of disinformation and fake news (2018)

[28] Moreno, Y., Nekovee, M., Pacheco, A.F.: Dynamics of rumor spreading in complex networks. Physical Review E **69**(6), 066130 (2004)

[29] Nekovee, M., Moreno, Y., Bianconi, G., Marsili, M.: Theory of rumour spreading in complex social networks. Physica A: Statistical Mechanics and its Applications **374**(1), 457 – 470 (2007). https://doi.org/10.1016/j.physa.2006.07.017

[30] Noymer, A.: The transmission and persistence of 'urban legends': Sociological application of age-structured epidemic models. Journal of Mathematical Sociology **25**(3), 299–323 (2001)

[31] P. Erdös, A.R.: On random graphs i. Publ. math. debrecen **6**(290-297), 18 (1959)

[32] Pastor-Satorras, R., Vespignani, A.: Immunization of complex networks. Physical review E **65**(3), 036104 (2002)

[33] Ratkiewicz, J., Conover, M.D., Meiss, M., Gonçalves, B., Flammini, A., Menczer, F.M.: Detecting and tracking political abuse in social media. In: Fifth international AAAI conference on weblogs and social media (2011)

[34] Ross, R.: An application of the theory of probabilities to the study of a priori pathometry.—part i. Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character **92**(638), 204–230 (1916)

[35] Ross, R., Hudson, H.P.: An application of the theory of probabilities to the study of a priori pathometry.—part ii. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **93**(650), 212–225 (1917)

[36] Shao, C., Ciampaglia, G.L., Varol, O., Yang, K.C., Flammini, A., Menczer, F.: The spread of low-credibility content by social bots. Nature communications **9**(1), 1–9 (2018)

[37] Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: A data mining perspective. ACM SIGKDD Explorations Newsletter **19**(1), 22–36 (2017)

[38] Tambuscio, M., Oliveira, D.F., Ciampaglia, G.L., Ruffo, G.: Network segregation in a model of misinformation and fact-checking. Journal of Computational Social Science **1**(2), 261–275 (2018)

[39] Tambuscio, M., Ruffo, G., Flammini, A., Menczer, F.: Fact-checking effect on viral hoaxes: A model of misinformation spread in social networks. In: Proceedings of the 24th International Conference on World Wide Web. p. 977–982. WWW '15 Companion, Association for Computing Machinery, New York, NY, USA (2015). https://doi.org/10.1145/2740908.2742572

[40] Tandoc Jr, E.C., Lim, Z.W., Ling, R.: Defining "fake news" a typology of scholarly definitions. Digital journalism **6**(2), 137–153 (2018)

[41] Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. Science **359**(6380), 1146–1151 (2018)

[42] Wang, J., Zhao, L., Huang, R.: Siraru rumor spreading model in complex networks. Physica A: Statistical Mechanics and its Applications **398**, 43–55 (2014)

[43] Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world'networks. nature **393**(6684), 440 (1998)

[44] Xiong, F., Liu, Y., jiang Zhang, Z., Zhu, J., Zhang, Y.: An information diffusion model based on retweeting mechanism for online social media. Physics Letters A **376**(30), 2103 – 2108 (2012). https://doi.org/10.1016/j.physleta.2012.05.021

[45] Zan, Y., Wu, J., Li, P., Yu, Q.: Sicr rumor spreading model in complex networks: Counterattack and self-resistance. Physica A: Statistical Mechanics and its Applications **405**, 159–170 (2014)

[46] Zhang, N., Huang, H., Su, B., Zhao, J., Zhang, B.: Dynamic 8-state icsar rumor propagation model considering official rumor refutation. Physica A: Statistical Mechanics and Its Applications **415**, 333–346 (2014)

[47] Zhao, L., Wang, J., Chen, Y., Wang, Q., Cheng, J., Cui, H.: Sihr rumor spreading model in social networks. Physica A: Statistical Mechanics and its Applications **391**(7), 2444 – 2453 (2012). https://doi.org/10.1016/j.physa.2011.12.008

[48] Zhao, L., Xie, W., Gao, H.O., Qiu, X., Wang, X., Zhang, S.: A rumor spreading model with variable forgetting rate. Physica A: Statistical Mechanics and its Applications **392**(23), 6146–6154 (2013)

[49] Zhao, Z., Zhao, J., Sano, Y., Levy, O., Takayasu, H., Takayasu, M., Li, D., Wu, J., Havlin, S.: Fake news propagates differently from real news even at early stages of spreading. EPJ Data Science **9**(1), 7 (2020)

Appendix A | Previous Work

Immunisation in Scale-free and Hierarchical Networks

Fábio Vital, Gonçalo Simões, Pedro Duarte - Network Science Project 2

December 2019

Introduction

In this second part of the project we extend our previous knowledge of random networks [1], with the Barabási–Albert (BA) and Dorogovtsev–Mendes–Samukin (DMS) models, and apply it to the field of epidemiology. We studied the spread of disease in these graphs and the result of different types of immunisation strategies. The Susceptible–Infectious–Recovered model was used and outbreak sizes (fraction of recovered nodes, f_R) were measured, for different viral strengths (β) with varying fractions of vaccination (f_V). We found indication that vaccination strategies which improve the chances of immunising a hub, for the same f_V, lead to smaller outbreaks. This is related to hubs being important infection vectors. Moreover we found that, although BA and DMS networks have similar outbreak sizes without any immunisation, the latter had lower f_R for all f_V, regardless of the vaccination method used, when compared to the BA networks. This indicates that topology plays a big role in disease spread, most likely, the local structure characteristics, resulting in a hierarchical network [1], improve the chances of immunisation strategies hitting a hub.

Simulations[6]

Networks. Simulations were performed using two previously studied models that produce scale-free networks [1]: Barabási–Albert (BA) and Dorogovtsev–Mendes–Samukin (DMS). Both models with multiple network sizes $N = 625, 1250, 2500, 5000, 10000$.

SIR. We used the same model as described on reference [2], where the disease spreading is assumed to be through static contact. The nodes of the graph represent people, and the edges their respective physical interactions. There are 4 states a node can be in — susceptible (S), infectious (I), recovered (R) or vaccinated (V). Susceptible nodes do not have the disease, but can get it. Infectious nodes have the disease and can spread it. Recovered nodes do not have the disease and can't get it. Vaccinated nodes don't have the disease and can't get it (immune). Disease outbreaks start at time t = 0, when a single node is randomly selected to be infected. β is the infection rate, it indicates how likely an infectious node spreads the disease to a susceptible neighbour. Infectious nodes recover at rate ν. In this setting, the infection and recovery times are independent exponential random variables, and an infectious node transfers a disease through an edge before getting recovered with probability $\beta/(\beta+\nu)$. The SIR model is controlled by these two variables independently. If we make $\nu = 1$ we'll

essentially be measuring in $1/\nu$ units of time. To simplify even further, we only perform an infection or a recovery on each iteration of the algorithm. With the probability that the next event being an infection given by

$$\frac{\beta M_{SI}}{\beta M_{SI} + N_I}, \tag{1}$$

where M_{SI} is the number of edges between susceptible and infectious nodes, and N_I is the number of infectious nodes (prevalence). If it's an infection event, a random M_{SI} edge is selected and the susceptible node becomes infectious. If there is no infection event, a recovery event takes place, where a random infectious node is selected to become recovered. Simulations were performed for values of $\beta = 1/16, 1/8, 1/4, 1/2, 1, 2, 4, 8, 16, 32$.

Number of samples. To obtain a good statistic, according to this study's objectives we performed simulations with several samples which were then averaged. The sample size used by P. Holmes and N. Litvak [2] is 300 000, which requires a substantial amount of computation time. In order to reach meaningful results and save on computation resources we experimented with sample sizes for a network of $N = 2500$, to check the standard deviation of the recovered nodes and time required. We chose a sample size of 300. On Figure 1a we observe that, for this amount of samples, the values are stable enough for both BA and DMS networks. Also, the time scales linearly with the number of samples, Figure 1b, making larger sample sizes too demanding, as we'd need to go above 500 samples to have statistical gains.

Vaccination methods and fractions. Before each SIR simulation, the generated BA and DMS graphs were vaccinated using one of 7 methods for fractions of $f_V = 0$, 0.1, 0.2,

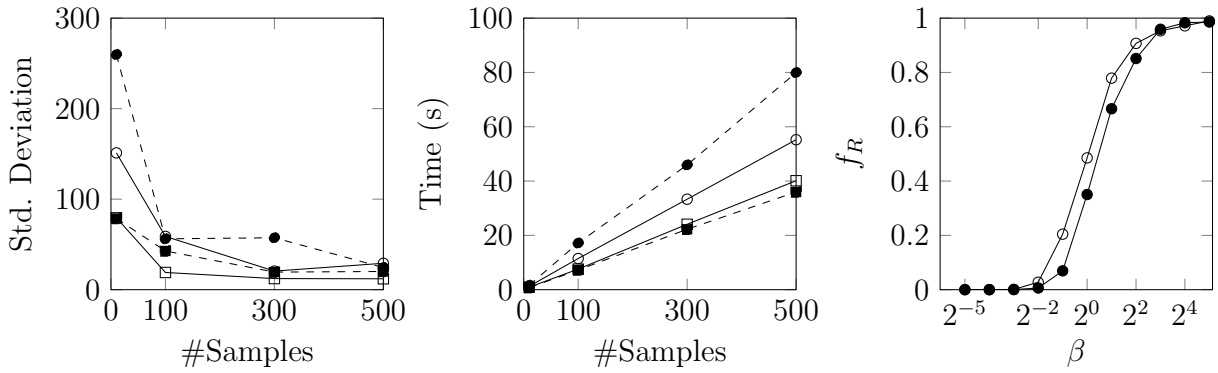

Figure 1: (a) Standard deviation of recovered nodes as a function of sample size, indicating low value variation above 100 samples, for a network of 2500 nodes – BA network with $\beta = 8$ (○) and $\beta = 16$ (●); DMS network with $\beta = 8$ (□) and $\beta = 16$ (■);
(b) Simulation time, in seconds, as a function of sample size for a network of 2500 nodes, BA network with $\beta = 8$ (○) and $\beta = 16$ (●); DMS network with $\beta = 8$ (□) and $\beta = 16$ (■);
(c) Fraction of recovered nodes as a function of β, showing that smaller values produce little to no outbreaks and rapidly scaling from $\beta = 1$ up to complete population outbreaks for values of $\beta > 4$. Plot data for BA (○) and DMS (●) networks with 10000 nodes.

0.3, 0.4, 0.5, 0.6, 0.7, 0.8, where 0 is the unvaccinated network.

- **Random.** Randomly select nodes and vaccinate them up to f_V.
- **DFS.** Depth First Search vaccination randomly selects a node, vaccinates it, and then vaccinates using a DFS search from that node, until f_V. If a node runs out of neighbours, before reaching f_V, a new origin node is randomly selected, repeating the process until f_V is reached. The process is repeated until f_V.
- **BFS.** Breadth First Search vaccination randomly selects a node, vaccinates it, then vaccinates all its neighbours, then the neighbours of the neighbours, iteratively until f_V. Like with DFS, when the nodes run out of neighbours, before reaching f_V, a new one is selected.
- **Bomb.** A BFS like method that vaccinates the random node and all its first order neighbours. Also selects a new random node until f_V.
- **Random walk.** Taken from [2], this method randomly selects a node as the origin, then traces a path, by randomly selecting neighbours, recursively, similar to the DFS. However, at each step, there is a chance (given by $\alpha/(k_i + \alpha)$, where k_i is the degree of the current node, $\alpha = 3$) of jumping from the path, by randomly selecting another origin. On this method we record the degree of all nodes that were visited and then select the ones with highest degree. To compare the same fractions among all methods, we record 10% extra nodes, above the needed f_V, and then choose from that set.
- **Acquaintance.** In this method we randomly select a node on the network and then vaccinate, also randomly, on of its neighbours. We repeat until we get f_V.
- **Hub.** This is a method which would have major costs in the real world, as determining the hubs in the network requires knowing all of its topology. It is a perfect and practically impossible scenario, but highly insightful in painting the picture of the roles of hubs in disease spread.

Results and Discussion

Network size. Despite having done simulations with several network sizes, we observe that the results don't appear to depend on this property. As seen on Figure 2, there are minimal finite size effects. For this reason we'll only show data plots for $N = 10000$ as they capture the full picture of the results obtained, for all network sizes.

Rate of infection (β). This parameter correlates with how fast the disease would spread in reality. The higher it is, the faster the network is infected. In Figure 1c, we see that f_R is directly proportional to β, which was expected given that the probability to infect increases with β, Eqn. 1. We can also observe that scale-free networks are very susceptible to infection spread, with both types having very similar responses to the different values of β. In a sense, the properties which confer them their name, allow disease to spread

uncontrolled, mainly the existence of hubs [3]. However, rates lower than 1, yield low f_R results in the current method of simulation, $f_R < 0.4$. Given that infectious nodes recover when no infection occurs, and the probability of infection is β, for lower values we tend to perform more recoveries than infections which stops spread. On the other end of the spectrum, for high rates ($\beta > 4$), the network is totally consumed by the disease.

Vaccination and topology. One of the main aims of epidemiology is to understand the role of vaccination and hence predict the level of vaccination necessary to eradicate a disease. As expected, vaccinating the network reduces f_R as prior infection routes become blocked completely or partially, requiring more steps to reach the different S nodes in the network, allowing for more opportunities of recovery. Once we introduce vaccination, regardless of the method that is used, we can observe a distinction between network types, BA and DMS, in how they behave. Although they have very similar responses without vaccination, as noted above, DMS shows a slight decrease in f_R as more of the network is vaccinated (higher f_V).

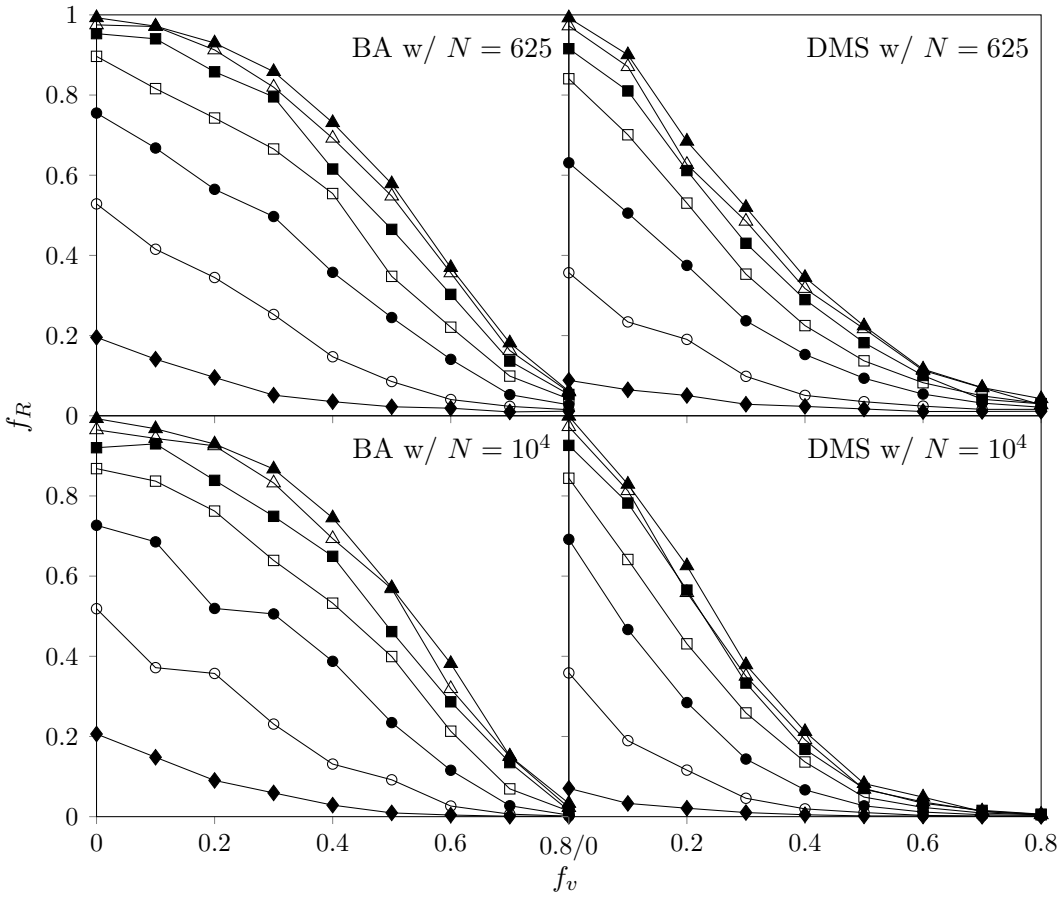

Figure 2: Fraction of recovered f_R as a function of vaccination fraction f_V for Random immunisation, showing an independence of results with the network size (little to no finite scale effects) and an already distinct response from both network types under immunisation, without specific strategy. Each plot has values for $\beta = 0.5$ (♦), $\beta = 1$ (○), $\beta = 2$ (●), $\beta = 4$ (□), $\beta = 8$ (■), $\beta = 16$ (△), $\beta = 32$ (▲).

This seems to indicate a greater resistance from the latter for disease spread, an apparent topological effect. From previous work, we know that the key differing factor of these two networks is their local clustering coefficient, with higher values for DMS following a power law [1]. A possible explanation is a higher probability of vaccinating key nodes, which prevents access to other levels of the hierarchy. This topology leads to the disease being in a sort of "trap", a sub-network, surrounded by vaccinated nodes, because the main paths of the hierarchy are now immune. Real world networks display the same power law relation for the clustering coefficient [3], which means that the network could be, in itself, an amplifier of the vaccination strategies. Others have observed this link between local structure and how it can be used to adjust immunisation strategies, improving response with less, but more targeted actions [5] in real networks like sexual HIV spreading in human sexual contact networks [3] or airline and school friendship networks for highly infectious diseases like SARS [5]. Although not using DMS networks, this effectively corroborates our observations on this topological characteristic's importance.

When we compare the different methods of vaccination, among the same network types, we see that randomly selecting nodes is the worst of the cases, with the exception of Random Walk on BA, where this only happens for $f_V \leq 0.5$.

Hubs. As was expected, vaccinating Hubs is very effective in limiting outbreak sizes, 3, indicating that finding them is key to blocking disease spread. It also puts a lower bound on how effective a method can be, for these types of networks in the context of the SIR model.

DFS and BFS (and Bomb). One interesting question we set out to answer was the effect of how we search the network for nodes to immunise. On 3 we see that DFS outperforms BFS in reducing outbreak sizes. This was expected due to the differences they have in the order in which they reach nodes from a given source. The first vaccinates along paths of increasing distance, and the second does the reverse, vaccinating all neighbours at a given distance, before stepping further. This makes DFS more likely to hit hubs, as it searches along paths, gaining more robustness to outbreaks. The Bomb method is worst than BFS because it fails to immunise to a greater distance, reducing its chances to find a hub, and exhausting vaccines in redundant nodes in the same neighbourhood. Furthermore, for DFS on DMS networks, an interesting effect can be observed for $f_V \geq 0.1$, with the values of f_R never reaching zero. This could be attributed, again, to an effect in the network's topology.

Acquaintance. Comparing the overall capacity to prevent disease spread, 3, this method yields the best results in both network types. On DMS networks this effect is reduced by the network's topology effect, but on BA there is a sharp difference with practically no spread for $f_V \geq 0.3$. This indicates that hub nodes are immunised very effectively, even more than with DFS. The difference can, most likely, be attributed to the fact that DFS "wastes" vaccines on every node along the path to a hub, while this method has more chances to hit a node, by playing more times the random selection. The random method also selects the same number of times, but it has to hit the hub. By selecting one of the neighbours, the Acquaintance method greatly increases its chances to hitting a hub (due to their high degree).

Random walk. This method has two performances, when compared to the Random vaccination: 1) for values $f_V \leq 0.5$ the outbreaks are smaller; 2) for higher values vaccinating more results in little, or no, gain. A possible explanation might be that the α value, which determines the probability to jump, is too low, and we end up walking along paths that always lead to low degree nodes. Furthermore, this method doesn't rollback like DFS, so

increasing the probability to select another node will make it more like Acquaintance, which yields the best chances of finding hubs.

Conclusions

It is clear that hubs are essential to disease spread and in that sense, also, key in preventing it. In the real world, if we consider a scale-free network of these types as a good approximation of the real network, our best bet is to use the Acquaintance method of immunisation, as it'll increase our chances of reducing outbreak sizes at lower costs. All other methods explore the topology in a reduced sub-set of the network, missing important paths of disease spread. Moreover, if there is a knowledge of the network, and high degree nodes can be identified, these should always be targeted for vaccination before the rest. In a sense this validates what actually happens in the real world, where medical staff and first responders for medical emergencies, are always the first to get immunisation. They can be seen as identifiable hubs. For the rest of the population the Acquaintance method should be the best response.

Perhaps to most important result is the role of topology when immunisation is carried. We saw that the DMS, and its hierarchical signature with high clustering coefficient, leads to improved performance of every vaccination strategy.

Future work. The most interesting avenue of future work would be the effect of topology on how effective vaccination methods are. This way, there can be two courses of action to combat disease spread. This could also be expanded to other conditions in the SIR model, as described below, and introducing dynamics into the network to see if the effect holds true when the network changes (assuming we can maintain the coefficient, otherwise it would be expected to result in what was observed for BA networks).

On the studied vaccination methods, it would be nice to see if a slight change in the DFS method produces any significant improvement. We vaccinated all nodes including leafs, but if we disregard a path, after finding a leaf, not vaccinating it, the effect, given a same f_V, should be improved, as more inner path nodes would be immunised. We could also combine it with Acquaintance, where the randomly selected node is not vaccinated, and determine if we could improve this strategy.

We focus on static contact model, which we could change in several ways to direct our studies to different aspects. Randomly infecting nodes up to a given distance, simulating air/water infection vectors; or introducing new states, such as death or incubation, modelling different kinds of diseases and how they affect their own success when hosts may die before infecting others. Targeted attacks, preferentially infecting leaf nodes or hubs, as patient zero, seeing if topology has the same results.

References

[1] F. Vital, G. Simões, P. Duarte, *Report – CRC Project 1 (Group 5)*, Ciência das Redes Complexas 2019–2020.

[2] Petter Holmes and Nelly Litvak, Cost-efficient vaccination protocols for network epidemiology. In PLOS Computational Biology, 2017, doi: 10.1371/journal.pcbi.1005696.

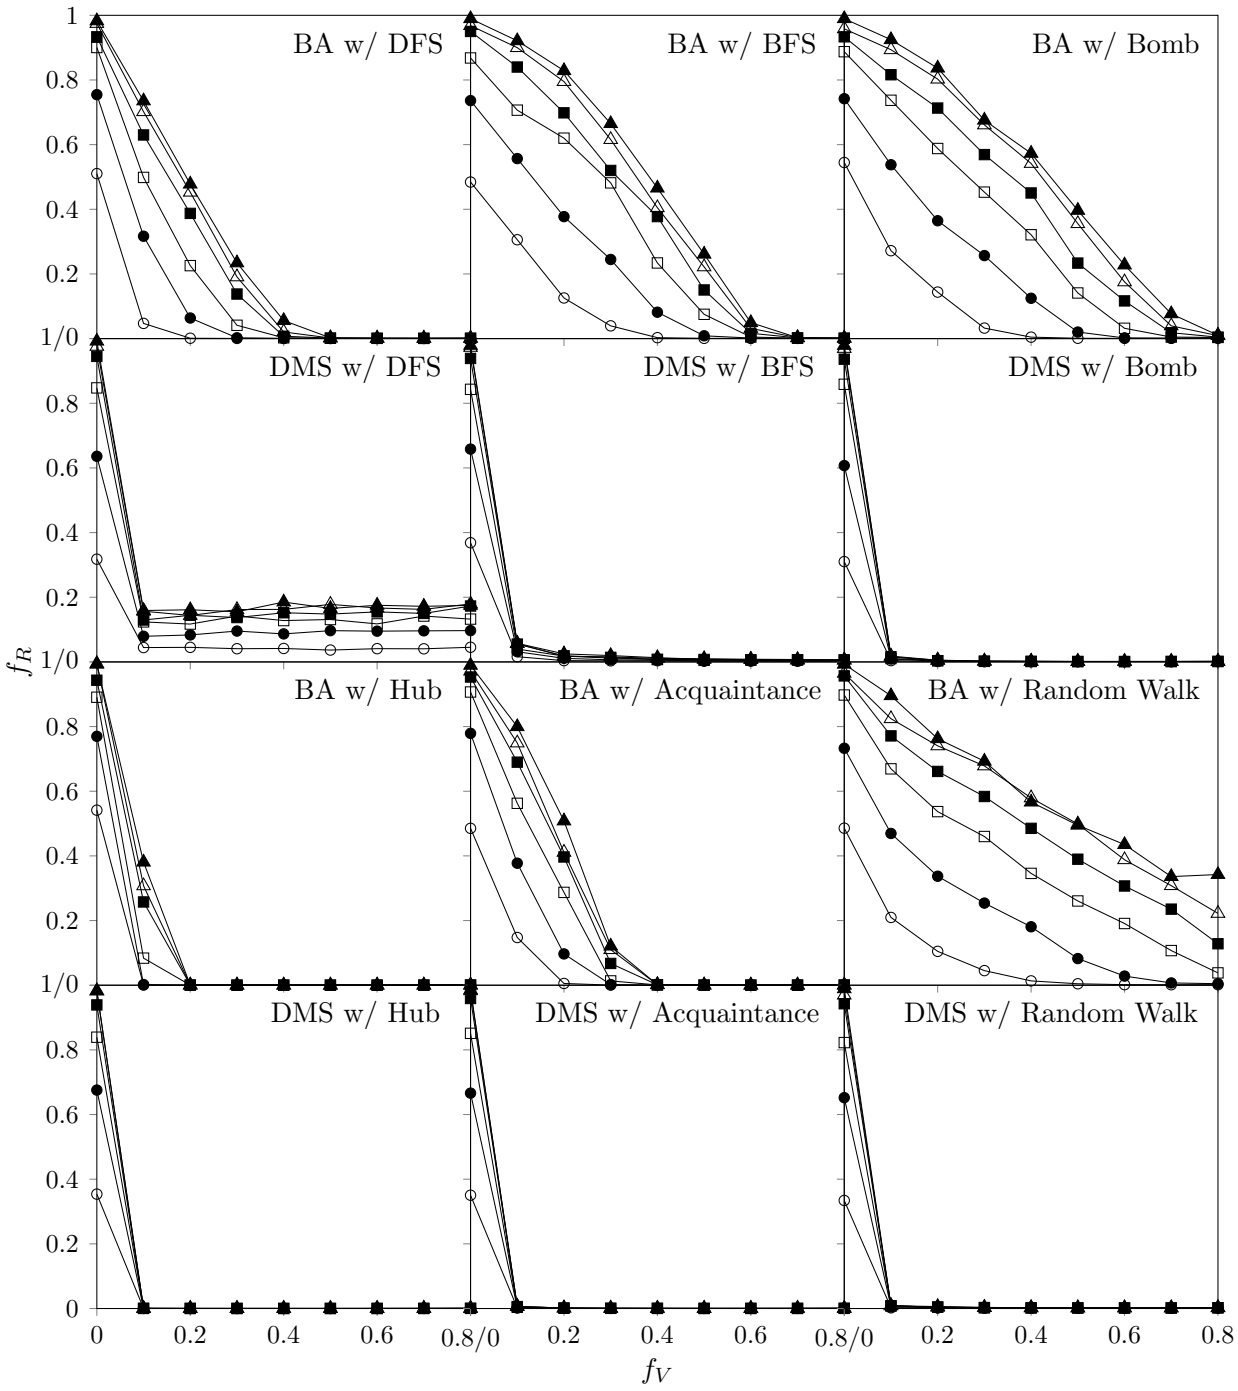

Figure 3: Recovered fraction f_R as a function of vaccinated fraction f_V showing the efficacy of the different active strategies (DFS, BFS, Bomb, Hub, Acquaintance and Random Walk) in reducing outbreaks. The DMS boosting effect is clearly seen on all plots, levelling their outcomes and suggesting that strategies are just options to the same end goal on networks with high clustering coefficient. Each plot has values for $\beta = 1$ (○), $\beta = 2$ (●), $\beta = 4$ (□), $\beta = 8$ (■), $\beta = 16$ (△), $\beta = 32$ (▲); results for all network types and all vaccination methods, except Random vaccination.

[3] Albert-Barabási, et al., Scale-Free and Hierarchical Structures in Complex Networks. In American Institute of Physics, 2003, doi: 10.1063/1.1571285.

[4] M. J. Keeling, The effects of local spatial structure on epidemiological invasions. In The Royal Society, 1999.

[5] Christian M. Schneider, et al., Suppressing epidemics with a limited amount of immunization units. In American Physical Society, 2011, doi: 10.1103/PhysRevE.84.061911.

[6] GitHub repository with the source code used for the simulations, Folder *proj2/*, at https://github.com/ghosw/crc-project.

Appendix B | Model Pseudocode Implementation

Algorithm 1 SIRV Implementation

```
1: class SIRV
2:     procedure CONSTRUCTOR(β, γ, δ, θ, ζ, timescale)
3:         states ← list("S","I","R","V")                          ▷ List of model's state names
4:         variables ← dictionary of model's variable names and their values (β, γ, δ, θ, ζ, timescale)
5:         transitions ← dict(
6:             "IS"  ← INFECT                                      ▷ Spreader infecting a susceptible
7:             "II"  ← STIFLING                                    ▷ Spreader being stifled by another spreader
8:             "IR"  ← STIFLING                                    ▷ Spreader being stifled by a stifler
9:             "IV"  ← VACCINATION STIFLING                        ▷ Spreader being stifled by a vaccinated
10:            "VS"  ← VACCINATE                                   ▷ Vaccinated vaccinating a susceptible
11:            "I"   ← FORGET                                      ▷ Spreader stifling himself via forgetting
12:        )
13:        initiate_model(states, variables, transitions)
14:
15:    procedure END CONDITION(network)
16:        n_infected ← number of infected states in the network
17:        if n_infected = 0 then return true
18:        return false
19:
20:    procedure STATES TO UPDATE(network)
21:        timescale ← model's timescale value
22:        if random_uniform_value(0, timescale + 1) <= 1 then return list("I")
23:        return list("V")
24:
25:    procedure INFECT(initiator, initiator_state, receiver, receiver_state, network)
26:        β ← model's β value
27:        if random_uniform_value(0, 1) <= β then return list(initiator_state, "I")
28:        return list(initiator_state, "V")
29:
30:    procedure STIFLING(initiator, initiator_state, receiver, receiver_state, network)
31:        γ ← model's γ value
32:        if random_uniform_value(0, 1) <= γ then return list("R", receiver_state)
33:        return list(initiator_state, receiver_state)
34:
35:    procedure VACCINATION STIFLING(initiator, initiator_state, receiver, receiver_state, network)
36:        ζ ← model's ζ value
37:        if random_uniform_value(0, 1) <= ζ then return list("R", receiver_state)
38:        return list(initiator_state, receiver_state)
39:
40:    procedure VACCINATE(initiator, initiator_state, receiver, receiver_state, network)
41:        θ ← model's θ value
42:        if random_uniform_value(0, 1) <= θ then return list(initiator_state, "V")
43:        return list(initiator_state, receiver_state)
44:
45:    procedure FORGET(node, state, network)
46:        δ ← model's δ value
47:        if random_uniform_value(0, 1) <= δ then return "R"
48:        return state
```