UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Developing Intervention Campaigns to Combat Anti-vaccine Opinion Propagation and their Impact on Disease Spread: A Computational Approach

by

Sarah Hamed Alahmadi

ORCiD: 0000-0001-9132-0779

A thesis for the degree of
Doctor of Philosophy

March 2026

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Doctor of Philosophy

Developing Intervention Campaigns to Combat Anti-vaccine Opinion Propagation and their Impact on Disease Spread: A Computational Approach

by Sarah Hamed Alahmadi

Information regarding vaccines from sources such as health services, media, and social networks can significantly shape vaccination decisions. In particular, the dissemination of negative information can contribute to vaccine hesitancy, thereby exacerbating infectious disease outbreaks. This study investigates strategies to mitigate anti-vaccine social contagion through effective counter-campaigns that disseminate positive vaccine information and encourage vaccine uptake, aiming to reduce the size of epidemics. In a coupled agent-based model that consists of opinion and disease diffusion processes, we develop different intervention strategies to mitigate the spread of vaccine misinformation subject to budget constraints. We first explore and compare different heuristics to design positive campaigns based on the network structure and local presence of negative vaccine attitudes. We examine two campaigning regimes: a static regime with a fixed set of targets, and a dynamic regime in which targets can be updated over time. We demonstrate that strategic targeting and engagement with the dynamics of anti-vaccine influence diffusion in the network can effectively mitigate the spread of anti-vaccine sentiment, thereby reducing the epidemic size. However, the effectiveness of the campaigns differs across different targeting strategies and is impacted by a range of factors. We find that the primary advantage of static campaigns lies in their capacity to act as an obstacle, preventing the large-scale clustering of emerging anti-vaccine communities, thereby resulting in smaller and unconnected anti-vaccine groups. On the other hand, dynamic campaigns reach a broader segment of the population and adapt to the evolution of anti-vaccine diffusion, not only protecting susceptible agents from negative influence but also fostering positive influence propagation within clusters of negative opinions.

We then frame this as a budget-constrained optimization problem, developing optimization-based strategies. Specifically, we explore two optimization objectives: minimizing the number of anti-vaccine opinion adopters and maximizing the number of pro-vaccine opinion adopters. We demonstrate that minimizing anti-vaccine opinion adoption is a more efficient strategy, as it promotes vaccine uptake within negative regions and thus blocks the diffusion of anti-vaccine influence. In contrast, the maximization approach operates near positive areas in the network, neither helping contain negative influence nor spreading positive influence throughout the network, resulting in poor control over the spread of anti-vaccine opinions and the corresponding epidemic size. Furthermore, we find that control performance deteriorates when interventions are delayed. This is due to the early formation of anti-vaccine communities that continue to grow, making them increasingly difficult to contain. We also introduce cluster-based optimization interventions aimed at curbing the expansion of existing communities. These campaigns identify and target critical regions in the network, effectively limiting the growth of the communities and thus reducing epidemic size.

Next, we systematically study the influence of network topology on the effectiveness of control strategies. We find that the efficiency of campaigns depends on both the network structure and the timing of the intervention. For early intervention, we demonstrate that minimizing anti-vaccine opinion adoption is an efficient strategy, particularly when incorporating detailed knowledge of individual's neighbourhoods, which helps improve the campaign by targeting agents who can positively influence their hesitant peers. For late intervention, shielding critical regions is an effective strategy in small-worlds and lattices. However, it is less effective in scale-free and random networks due to the distinct patterns of cluster formation in these networks which tend to produce large clusters from the early stages of their emergence. In addition, although controlling negative diffusion becomes more challenging the longer the intervention is delayed, we find that it can be managed more efficiently and with fewer resources in small-world and regular lattice networks than in scale-free and random networks.

Last, we study interventions in multi-layer networks considering the different venues of social interactions. We find that, virtual-based interventions generally outperform physical-based interventions across most campaigns. An exception is the cluster-based campaigns, where controlling physical anti-vaccine communities performs better than controlling virtual anti-vaccine communities. In addition, missing social contact data from the virtual layer can degrade control performance, particularly when the layer has a scale-free structure. Furthermore, campaign behavior in multi-layer networks is consistent with that in single-layer networks, suggesting that the choice of intervention strategy should align with the structure of the dominant information layer.

In conclusion, we make several practical recommendations on targeted interventions to counter the diffusion of vaccine misinformation and its impact on public health. We provide insights into effective implementation for policymakers and public health administrators, while also highlighting the limitations and operational challenges that may affect the effectiveness of these interventions.

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as: Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Modelling the mitigation of anti-vaccine opinion propagation to suppress epidemic spread: A computational approach. *PLOS ONE*, 20(3):e0318544, 2025b
   Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread. In Hocine Cherifi, Murat Donduran, Luis M. Rocha, Chantal Cherifi, and Onur Varol, editors, *Complex Networks & Their Applications XIII*, pages 87–99, Cham, 2025a. Springer Nature Switzerland. ISBN 978-3-031-82439-5
   Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread: Impact of network structure. *Applied Network Science*, 10(1):59, 2025c

Signed:.. Date:..................

Acknowledgements

First and foremost, all praise and thanks are due to Allah, who granted me strength, patience, and guidance throughout this journey. This work would not have been possible without His grace and blessings.

I would like to express my deepest gratitude to my supervisors, Dr. Markus Brede, Prof. Rebecca Hoyle, and Dr. Michael Head, for their unwavering support, insightful feedback, generous time, and continuous encouragement. Your invaluable guidance has been instrumental in shaping this research and helping me grow as a researcher. I am truly honoured to have had your mentorship.

I extend my deepest gratitude to my family. To my dear parents and family—thank you for your unwavering belief and support, even from afar. To my beloved husband, Mohammad—thank you for your constant support and encouragement; they have been my source of strength during this journey. To my wonderful children, Lamar, Abdulaziz, and Hussam—you are my constant source of inspiration and motivation. Your love, sweet encouragement, and innocent prayers gave me strength and hope when I needed them most.

I also extend my sincere gratitude to my friends for their kindness, unwavering support, and the beautiful moments we shared throughout this journey.

To all my colleagues and friends in Southampton — thank you all.

Chapter 1

Introduction

Throughout history, diseases have posed a constant threat to human health and wellbeing, with the Covid-19 pandemic serving as a recent and striking example. Vaccination is a vital tool for combating disease prevalence. However, in many instances, the availability of vaccines does not necessarily lead to their uptake owing to several factors, including vaccine hesitancy. Vaccine hesitancy, defined by the WHO as the reluctance or refusal to get vaccinated despite the availability of vaccines, poses a significant challenge to public health initiatives aimed at preventing outbreaks of vaccine-preventable diseases. It has been identified by the WHO as one of the top ten threats to public health (WHO, 2019). One significant factor contributing to it is exposure to vaccine misinformation (Dubé et al., 2013; Burki, 2019; Hotez et al., 2020; Pierri et al., 2022; Lee et al., 2022).

Information propagation in social networks can significantly influence overall public opinion and behavior. With the advent of social media, this process has become faster, more complex, and far-reaching. In particular, the spread of health-related misinformation plays a significant role in shaping people's willingness to receive vaccines or adopt precautionary measures (Vinck et al., 2019; Hotez et al., 2020; Lee et al., 2022). Such negative behaviors can spread through social interactions, potentially undermining public health efforts to contain epidemics. This coupled dynamic leads to the so-called Infodemic, defined by the WHO as an overabundance of information, including false or misleading content, causing confusion, harmful behaviors, mistrust in health authorities, and potentially intensifying or lengthening disease outbreaks (World Health Organization). This phenomenon presents a serious challenge to public health and underscores the need to manage vaccine misinformation, due to their potential to alter the trajectory of epidemics.

Several historical outbreaks exemplify the detrimental impact of misinformation on public health efforts (Vinck et al., 2019; Hotez et al., 2020; Lee et al., 2022). The spread of misinformation during the Covid-19 pandemic resulted in increased vaccine

hesitancy and low vaccination rates (Lee et al., 2022; Pierri et al., 2022). Similarly, during the Ebola outbreak in North Kivu in the DR Congo, a notable association was observed between the mistrust in health systems and the belief in misinformation, which led to a reduced inclination to adopt preventive behaviors, including acceptance of Ebola vaccines (Vinck et al., 2019). Additionally, the year 2019 witnessed a resurgence of measles due to a significant decline in vaccine coverage, resulting from the spread of anti-vaccine attitudes as a significant driver (Hotez et al., 2020). The year 2023 witnessed another resurgence of measles cases, attributed to lower vaccination uptake (WHO, 2024). This disease remains an ongoing global threat, with outbreaks continuing to emerge in various regions (Do and Mulholland, 2025). Such examples highlight the pivotal role of accurate information in managing and mitigating the effects of misinformation associated with various infectious diseases.

The intertwined dynamics between information and disease have attracted increasing attention in the literature through information–disease coupled models, addressing the extent to which information can influence the scale of epidemics. Information plays a dual role in public health and disease control, with the potential for both negative and positive effects. Specifically, information can suppress disease spread by triggering self-protective behaviors through the dissemination of positive information, such as raising awareness through disease-related information, which can potentially reduce the size of an epidemic (Ruan et al., 2012; Funk et al., 2009; da Silva et al., 2019; Sontag et al., 2022). Conversely, it can exacerbate disease spread, as in the case of negative information that acts as a stimulant for negative behavioral responses, such as the spread of vaccine misinformation over a social network which negatively influences people's inclination to vaccinate (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013; Dorso et al., 2017; Dai et al., 2019; Mehta and Rosenberg, 2020; Mumtaz et al., 2022; Chen et al., 2022b; Yin et al., 2022).

Opinions regarding vaccines spread among individuals as a social contagion, significantly shaping vaccination coverage and, consequently, the severity of outbreaks. In particular, anti-vaccine opinion adopters tend to cluster in social networks (Salathé and Khandelwal, 2011; Dubé et al., 2015; Omer et al., 2009; Yuan et al., 2019). This leads to the formation of communities consisting of unprotected individuals, representing vulnerable zones for disease transmission. Such communities pose a threat to public health and prevent attainment of high vaccination rates (Lieu et al., 2015) as well as impeding the effects of full herd immunity and high vaccination rates (Salathé and Bonhoeffer, 2008; Gromis and Liu, 2022). Numerous studies have investigated the risk of anti-vaccine clusters, focusing for instance on, the relationship between anti-vaccine communities and the epidemic size (Salathé and Bonhoeffer, 2008; Salathé and Khandelwal, 2011; Campbell and Salathé, 2013; Dorso et al., 2017), correlations between spatial clusters and the associated vaccination rates (Lieu et al., 2015), and using social media data to identify vaccine opinion clusters and

vaccination rates (Bello-Orgaz et al., 2017; Salathé and Khandelwal, 2011; Cheng et al., 2023). Studying the diffusion of misinformation and vaccine opinions not only helps in understanding the range of influences on people's attitudes but also helps in understanding the flow of attitudes across the social network and, therefore, allows for more effective intervention and control strategies.

Despite the significant attention devoted to understanding the coupled dynamics of information and epidemics, there have been limited efforts to explore the impact of targeted external interventions aimed at mitigating the diffusion of vaccine misinformation on the spread of infectious diseases (Dorso et al., 2017; Ancona et al., 2022; Yin et al., 2023; Fügenschuh and Fu, 2023). Further research is needed to design effective intervention strategies to counter the spread of anti-vaccine information within coupled diffusion processes and capture its impact on disease prevalence. Scholars have thoroughly explored strategic targeting in other fields, such as vaccination campaigns (Prieto Curiel and González Ramírez, 2021), and in single diffusion processes related to the influence maximization problem (Romero Moreno et al., 2021, 2020; Erkol et al., 2019) and misinformation mitigation problem (Budak et al., 2011; Zhang et al., 2015; Yang et al., 2019a). However, the development of such targeting strategies considering the complex interplay between information diffusion and disease spread as coupled processes remains limited.

In this thesis, we explore methods to combat the spread of anti-vaccine opinions through intervention campaigns that disseminate positive vaccine-related information, aiming to effectively mitigate the spread of anti-vaccine influence and amplify the spread of positive influence, with the ultimate goal of suppressing epidemic spread. Empirical research highlights the significance of healthcare providers' advice in overcoming vaccination hesitancy and promoting vaccine uptake through the dissemination of accurate information about vaccines (Adeyanju et al., 2021; Marzo et al., 2022). In addition, the effectiveness of many real-world interventions in addressing vaccine hesitancy and promoting vaccination uptake has been demonstrated through a variety of campaigns. For example, Covid-19 vaccine uptake improved in areas targeted by campaigns aimed at mitigating vaccine hesitancy in Nigeria (Evans et al., 2023) and among African immigrants in Philadelphia (Koku et al., 2025). Similarly, infant vaccination rates have been enhanced through campaigns targeting mothers during postpartum hospitalization (Lemaitre et al., 2019).

Therefore, through a coupled agent-based model that integrates opinion and disease spread processes, this study investigates how different methods to control the diffusion of anti-vaccine opinions influence epidemic dynamics. It aims to advance our theoretical understanding of the co-evolving dynamics between individuals' behavior and disease, as well as the role of external influence in shaping these dynamics. This deeper understanding enables the identification of key factors that

affect the effectiveness of different control strategies in curbing the propagation of negative influence and achieving the primary objective of suppressing epidemic spread. It further allows the development of effective interventions tailored to different scenarios. To this end, several research gaps have been identified in the field and remain to be addressed. In the following, we detail these gaps and key questions, as well as our contributions to addressing them.

1.1 Research Gaps and Questions

This research explores how behavioral responses to vaccination influence the dynamics of disease spread. It addresses the problem of mitigating the spread of vaccine misinformation in social networks through targeted intervention strategies that promote positive vaccine information, with the primary objective of containing the transmission of infectious diseases.

The intertwined dynamics of diseases and information propagation are complex, as in reality they can spread through different social networks with distinct structural properties. In real social systems, individuals maintain a variety of relationships, such as with family, friends, colleagues, or others, both online and offline. These different types of interactions often represent distinct contact structures and correspond to different spreading processes such as information and diseases. This diversity and complexity can be modeled using multi-layer networks in which different types of relationships are captured as separate network layers. However, directly addressing the problem of mitigating misinformation propagation and its impact on disease dynamics in multi-layer networks is complex as it requires a systematic understanding of how different factors interact to shape their coupled dynamics. To address this complexity, we investigate the problem progressively by decomposing it into subproblems addressing distinct questions. In particular, we begin with single-layer networks, which allows us to first design different intervention strategies and evaluate their effectiveness across various network topologies. Then, we extend our investigation to more realistic settings that incorporate diverse types of social interactions through multi-layer networks.

In the following, we enumerate the open questions and challenges identified in the field and addressed in this thesis:

1. **Impact of different intervention strategies to combat the propagation of vaccine misinformation on disease dynamics:** one major gap identified in the literature is the limited research exploring the impact of mitigating the spread of vaccine misinformation through external intervention campaigns on epidemic spread (Dorso et al., 2017; Yin et al., 2023; Fügenschuh and Fu, 2023). Strategic

interventions to control information diffusion and opinion adoption have been widely studied in the literature as separate problems, known as influence maximization (Romero Moreno et al., 2021, 2020) and misinformation mitigation (Budak et al., 2011; Liu et al., 2016; Tong et al., 2018; Yang et al., 2019a). Inspired by these approaches, leveraging control techniques to strategically combat the diffusion of vaccine misinformation has the potential to manage epidemic outbreaks; however, this area remains largely unexplored. To fill this knowledge gap, we study how different strategies aimed at limiting the spread of vaccine misinformation can contribute to curbing the epidemic. Specifically, we address the mitigation of the diffusion of anti-vaccine opinions as a misinformation mitigation problem, where the primary objective is to seek a set of agents to seed positive influence in a social network in a way that amplifies its reach and combats the propagation of anti-vaccine opinion diffusion. Targeting methods can be broadly categorized into heuristic approaches, such as selecting central nodes in the network, and optimization-based approaches using optimization techniques, such as the gradient ascent method for optimal resource allocation (Romero Moreno et al., 2021). Additionally, there are two methods for the seed set selection: static and dynamic. In the static approach, which has traditionally been used (Budak et al., 2011; Yang et al., 2019a; Fügenschuh and Fu, 2023), the seed set is chosen at the beginning of the campaign launch and remains unchanged throughout the campaign. In contrast, in the dynamic approach, the seed set is selected repeatedly in different rounds based on specific criteria (Wijayanto and Murata, 2019; Shi et al., 2019). This gap raises the following questions:

(a) To what extent can different heuristic methods to seed pro-vaccine influence mitigate the spread of anti-vaccine opinions, and how does this influence the final epidemic size?

(b) How can we design adaptive targeting strategies that account for the evolving dynamics of negative opinion propagation?

(c) Which seed selection method is more effective, static or dynamic, in countering anti-vaccine opinion spread? and why?

(d) How do different optimization objectives, namely minimizing the number of individuals adopting anti-vaccine opinions and maximizing the number of individuals adopting pro-vaccine opinions, contribute to countering the diffusion of anti-vaccine influence within the network, and what is their impact on the resulting epidemic size?

2. **Risk associated with anti-vaccine communities:** empirical research highlights both the geographic clustering of unvaccinated individuals (Dubé et al., 2015) and the formation of distinct anti-vaccine clusters in online social networks as a result of vaccine misinformation diffusion (Salathé and Khandelwal, 2011).

Several studies demonstrate that the emergence of anti-vaccine communities correlates with higher extent of disease spread (Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Gromis and Liu, 2022). They also undermine herd immunity and create vulnerable zones for disease transmission (Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Gromis and Liu, 2022). Understanding the evolution of these emerging communities is therefore critical for effectively containing disease spread. Given the significant risk posed by the emergence of these communities and the limited research addressing this problem, we explore the following questions:

(a) How does the initial size of existing anti-vaccine communities before intervention influence the effectiveness of containment strategies? What is the impact of the timing of intervention on control performance? Which campaign performs best, which performs worst, and what are the underlying reasons for these outcomes?

(b) How can we design a strategy to mitigate the interconnectivity and expansion of anti-vaccine clusters and prevent the emergence of large-scale communities? To what extent does this contain negative diffusion and reduce the resulting epidemic size?

3. **Impact of network topology on intervention performance:** given that networks exhibit considerable structural variability (Boccaletti et al., 2006; Newman, 2018), which influences how information, behaviors, or pathogens spread (Hein et al., 2006; Varga, 2017; Pérez-Ortiz et al., 2022), they can therefore significantly shape opinion propagation and evolution. It is essential to explore how network topology affects the efficacy of various interventions, and whether strategies that perform well in one type of network will maintain their effectiveness in others. A robust understanding of how anti-vaccine opinions propagate and anti-vaccine communities emerge in different network topologies is crucial for assessing their impact on the proposed intervention strategies. Therefore, we extend our work and investigate the following questions:

   (a) How does network topology influence system dynamics and the effectiveness of intervention strategies?

   (b) What are the structures and formation patterns of anti-vaccine communities across different network topologies?

4. **Capturing complex interaction behavior through multi-layer networks:** the different venues of contagion processes, namely information and disease, and their varying network structures have recently attracted researchers' attention. In particular, studies have explored how the spread of information and opinion adoption through different social interaction platforms, such as social media, can influence the physical transmission of disease in face-to-face networks using

multi-layer network frameworks (Dai et al., 2019; Yin et al., 2022; Huang et al., 2021; Chen et al., 2022b; Fügenschuh and Fu, 2023). In such frameworks, each layer captures a different type of social relationship among agents, with one layer representing information diffusion and another representing physical contact through which disease spreads. Therefore, considering the different network structures underlying the various spreading processes is crucial for understanding their interplay. In this regard, critical questions arise:

(a) How do system dynamics behave when information diffuses through a network structure distinct from the disease transmission network? What structural patterns emerge in anti-vaccine communities under these conditions? How do campaigns perform compared to those in single-layer control scenarios?

(b) To what extent does missing contact data from the virtual social network undermine the effectiveness of control strategies? Most importantly, which network layer plays the most critical role in achieving effective control - the virtual layer or the physical layer?

This thesis addresses the research questions outlined above, and the following section presents the main contributions made to address them.

1.2 Research Contributions

As mentioned above, this thesis explores the effects of external influence on the dynamics of vaccine-related opinions and how this, in turn, affects the dynamics of disease spread. To this end, it focuses on three key pillars. First, an important part of this thesis involves developing intervention strategies to seed pro-vaccine influence in a network to mitigate the diffusion of anti-vaccine influence. In this context, we also examine the impact of different factors, including budget constraints, strength of social influence, data uncertainty, and the timing of interventions, on the performance of these strategies. Second, we analyze the formation and evolution of anti-vaccine communities, exploring how to effectively contain their growth. Third, we study the impact of network topology on system dynamics, as well as how information propagation in different social networks affects disease dynamics in the physical network in a multi-layer framework. Through these investigations, our aim is to develop a theoretical understanding of the coupled dynamics of social behavior and disease, and to uncover the complex interplay between these dynamics and external interventions.

These three pillars correspond to the following five contributions, each addressed in a separate chapter as outlined below. The first four contributions investigate the

problem in single-layer networks, while the fifth one relates to multi-layer networks. The contributions are detailed as follows:

1. First, we develop a computational framework to examine the effects of mitigating anti-vaccine opinions on the size of epidemics. This model integrates external intervention campaigns, pro- and anti-vaccine opinion diffusion, and disease spread. Second, we explore and compare different heuristics aimed at identifying a set of nodes within the network to seed positive influence, based on network topology and vaccine-related information. These heuristics include both static and dynamic control approaches with different selection criteria. The static control involves targeting a fixed set of individuals, whereas the dynamic control involves an adaptive campaign that responds to the evolving dynamics of the anti-vaccine propagation. We examine their control effectiveness in counteracting the diffusion of anti-vaccine opinions and, consequently, in containing the spread of disease. This addresses research questions 1a,1b, and 1c, as outlined in Section 1.1, with a detailed discussion in Chapter 3.

2. We introduce a novel approach by formulating the problem as an optimization problem and proposing optimization-based heuristics. Specifically, we establish a mathematical framework that describes opinion dynamics in our model. We then define two objective functions and formulate two optimization problems aimed at: (1) minimizing the number of anti-vaccine opinion adopters; (2) maximizing the number of pro-vaccine opinion adopters. We examine the impact of these objectives on controlling opinion diffusion and the resulting epidemic size. We also analyze their targeting mechanisms to understand factors that shape their behavior and performance. We then extend our investigation to evaluate control performance in scenarios where an ongoing spread of rumors and misinformation already exists and has influenced a portion of the population, reflecting more realistic settings. This addresses research questions 1d and 2a, as outlined in Section 1.1, with a detailed discussion in Chapter 4.

3. We next address the problem of mitigating the growth of existing anti-vaccine communities. Specifically, we propose novel cluster-based optimization campaigns to control the evolutionary dynamics of anti-vaccine communities. These strategies identify and target critical areas in the network to limit the expansion of these communities and prevent the formation of larger communities. They also consider the potential risk posed by each community by considering their current size, prioritizing the control of larger communities over smaller ones. This addresses research question 2b, as outlined in Section 1.1, with a detailed discussion in Chapter 5.

4. We provide a systematic investigation of how network topology impacts the performance of the proposed intervention campaigns, identifying which

campaign performs best in each topology and the underlying reasons for these outcomes. Additionally, we conduct a thorough analysis of the clustering patterns and the formation of anti-vaccine communities in each topology. This addresses research questions 3a and 3b, as outlined in Section 1.1, with a detailed discussion in Chapter 6.

5. After gaining insights and a deeper understanding of the performance of each intervention strategy across a variety of network topologies and the factors that influence their performance in single-layer networks, it is of interest to explore scenarios in which information spread occurs on a different network structure than that of disease transmission which can be captured using multi-layer networks. This reflects a more realistic setting by considering distinct social channels through which individuals exchange opinions. Therefore, we extend our work to a multi-layer model where opinion propagation occurs in one layer and disease transmission in another. In this framework, we conduct a detailed analysis considering several key factors including network structures, the correlation between layers, and which layer to target. Specifically, we examine whether intervention through virtual networks or intervention through physical (face-to-face) networks are more effective. This addresses research questions 4a and 4b, as outlined in Section 1.1, with a detailed discussion provided in Chapter 7.

1.3 Related Publications and Contributions

This section outlines the publications and other contributions relevant to this thesis. The publications resulting from this research are listed below:

1. **Alahmadi, S., Hoyle, R., Head, M., & Brede, M. (2025). Modelling the mitigation of anti-vaccine opinion propagation to suppress epidemic spread: A computational approach. PLOS ONE, 20(3), e0318544.** (Published)
2. **Alahmadi, S., Hoyle, R., Head, M., & Brede, M. (2024). Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread. In Complex networks & their applications XIII: Volume 4, proceedings of The Thirteenth International Conference on Complex Networks and their Applications COMPLEX NETWORKS 2024 (pp. 87–99). Springer Nature Switzerland.** (Published)
3. **Alahmadi, S., Hoyle, R., Head, M., & Brede, M. (2025). Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread: Impact of network structure. Applied Network Science, 10(1), 59.** (Published)

Additional contributions, including posters and talks, are outlined below:

1. Poster presented at the Dynamic Days US conference (DDUS), 2023.
2. Extended abstract presented at French Regional Conference on Complex Systems (FRCCS), 2023.
3. Talk delivered at the British Applied Mathematics Colloquium (BAMC), 2023.
4. Extended abstract presented at the 15th International Conference on Complex Networks, 2024.
5. Seminar given to JUNIPER (Joint UNIversities Pandemic and Epidemiological Research), 2024.

1.4 Thesis Structure

The structure of the thesis is as follows. Chapter 2 provides an in-depth exploration of fundamental concepts needed to understand the work, including complex networks, opinion diffusion models, epidemiological models, and a review of relevant studies in coupled information-disease models. Chapter 3 presents the model description and the proposed heuristic-based intervention campaigns, along with a comprehensive analysis of the performance of each campaign. Chapter 4 introduces optimization-based campaigns, particularly individual-based approaches aimed at minimizing the number of anti-vaccine opinion individuals and maximizing the number of pro-vaccine opinion individuals. Chapter 5 extends our discussion and introduces cluster-based optimization campaigns aimed at mitigating the formation of anti-vaccine communities. Chapter 6 investigates the influence of network topology on the effectiveness of the proposed campaigns. Chapter 7 advances the discussion by investigating control performance in multi-layer networks. Finally, Chapter 8 provides a summary of the main findings and conclusions of the study, as well as directions for future work.

Chapter 2

Background Theory and Literature Review

The dynamics of social systems and the interactions among different spreading processes, e.g., information, behaviors, or diseases, are inherently complex. Studying interventions and control in such intertwined systems requires a thorough understanding of modeling techniques that accurately capture their complexity, components, and diverse properties. Networks are powerful tools for modeling social systems, enabling an abstract representation of real-world settings while capturing the structure and connectivity of social ties.

The primary goal of this research is to explore methods for suppressing the spread of vaccine-preventable infectious diseases by mitigating the propagation of vaccine misinformation in a social network through effective network-based interventions that promote positive vaccine-related information. Specifically, we aim to develop intervention strategies to select target agents from the entire population in a way that effectively combats the spread of negative information. Achieving this goal requires an understanding of the fundamental modeling methods for key system components: social networks, diffusion processes, and principles of control theory.

In this chapter, we present the fundamental concepts essential for understanding the present study, including the modeling approaches used to model the different components of the system. Specifically, Section 2.1 provides a review of the basic concepts in network science. Section 2.2 describes the opinion diffusion methods used to model the spread of information and opinions. Next, Section 2.3 discusses methods for modeling epidemic spread and outlines the standard models commonly used in the field.

We then proceed to relevant studies on coupled models that integrate information and disease spreading processes, describing the interacting dynamics of these processes in

Section 2.4. In addition, we provide an overview of the fields of information and epidemic control in Section 2.5, reviewing existing techniques for managing the spread of information (Section 2.5.1), and highlighting existing work on controlling disease spread (Section 2.5.2).

2.1 Network Properties and Models

Social systems consist of interconnected individuals whose interactions form complex webs of relationships. These complex systems can be modeled as networks, where individuals are represented as nodes and their interactions as edges. Networks represent an efficient framework for modeling social systems, enabling systematic analysis of how social interactions govern system dynamics and influence collective behavior.

In network theory, a graph (or network) G consists of a set of vertices (V) and a set of edges (E), $G = (V,E)$. Each vertex $v_i \in V$ represents an individual, and each edge $e_{ij} \in E$ represents a contact (interaction) between individuals v_i and v_j (Boccaletti et al., 2006). A graph G can be represented by an $N \times N$ square matrix called the *adjacency matrix*, where N is the number of vertices. The entry of this matrix, a_{ij} $(i, j = 1, \ldots, N)$, is equal to one if an edge exists between nodes i and j, and zero otherwise. The nodes that are directly connected to a node i are referred to as its neighbors.

Real-world networks exhibit diverse topological structures, each characterized by distinct structural properties (Boccaletti et al., 2006; Newman, 2018). Social networks exemplify this diversity, as they can exhibit different structures in which different types of social relationships display distinct connectivity patterns. For example, many social networks exhibit small-world properties (Watts and Strogatz, 1998; Barrat and Weigt, 2000; Telesford et al., 2011) as well as scale-free structures (Barabási and Albert, 1999; Mislove et al., 2007; Sadikov and Martinez, 2009). To quantify these structural variations, several metrics are commonly used, including node degree, degree distribution, clustering coefficient, density, and others (Boccaletti et al., 2006; Newman, 2018). The degree k_i of node i is defined as the number of edges connected to the node, and can be calculated from the adjacency matrix as:

$$k_i = \sum_{j=1}^{N} a_{ij}. \tag{2.1}$$

Accordingly, the average degree, $\langle k \rangle$, of a network can be calculated as follows (Newman, 2018):

$$\langle k \rangle = \frac{1}{N} \sum_{i=1}^{N} k_i. \tag{2.2}$$

In addition, the density of a graph, D, is the ratio of the number of edges present to the maximum possible number of edges, as follows:

$$D = \frac{2m}{N(N-1)}, \tag{2.3}$$

where m represents the number of edges present.

A fundamental characteristic of a graph G is its *degree distribution* P_k, defined as the fraction of nodes in the graph G with degree k. The degree distribution provides insight to understand the connectivity patterns within a network. For example, a skewed degree distribution indicates the presence of hub nodes, which are few nodes with a high number of connections while the majority of nodes have few connections.

Another characteristic of a graph G is the clustering (transitivity) which represents a standard feature of acquaintance networks, in which two friends of a node are likely to know each other (Boccaletti et al., 2006). This feature measures the extent to which nodes in a graph are clustered together. In graph representation, this type of connection is referred to as triangles, and transitivity means a high presence of triangles. The clustering of a network can be measured by the average clustering coefficient CC as follows:

$$CC = \frac{1}{N}\sum_{i=1}^{N} c_i, \tag{2.4}$$

where c_i is the local clustering coefficient for node i and is calculated as follows:

$$c_i = \frac{2e_i}{k_i(k_i-1)}, \tag{2.5}$$

where e_i is equal to the number of edges between neighbors of node i (Boccaletti et al., 2006).

Another important set of structural metrics relates to network navigation, including path, diameter, and average shortest path length (Boccaletti et al., 2006; Newman, 2018; Bianconi, 2018a). A path between two nodes is defined as a sequence of connected nodes, with each node linked to the subsequent node by an edge. Accordingly, the path length is the number of edges traversed between two vertices. A geodesic path is the shortest path between a pair of vertices. From these metrics, two important global measures can be derived. First, the diameter of a network, defined as the longest shortest path in the network. Second, the average shortest path length, defined as the average of all shortest paths between any two nodes of the network.

Beyond network definitions and fundamental metrics, capturing multiple types of social interactions and their distinct structural patterns is an important dimension in modeling real-world social networks. In real social systems, people have different

kinds of relationships, such as with family, friends, colleagues, and others. Some might be virtual relationships, while others might be face-to-face interactions. These different types of interactions often represent distinct contact structures and correspond to different spreading processes. For example, information can spread through any type of interaction, whether virtual or in-person, whereas disease transmission typically requires physical proximity and cannot occur through virtual contact.

To model these diverse interactions, a single-layer network, where all interactions are aggregated to a single network, may not accurately reflect the complexity of real-world systems. To address this limitation, multi-layer networks have been used, in which different types of relationships are modeled as separate network layers with interconnections between them (Bianconi, 2018b; Aleta and Moreno, 2019). Such frameworks enable more accurate modeling of complex systems and are able to preserve the distinct characteristics associated with different types of interactions or spreading processes. We will discuss this in more detail in section 2.1.2.

Several models have been proposed to generate synthetic networks that exhibit properties similar to real-world social networks. These models enable control over structural properties, allowing the study of a variety of network structures to better understand their impact on system dynamics. Beyond basic network metrics, real-world networks often exhibit distinct topological properties, such as the small-world effect and preferential attachment (power-law degree distributions). In the following, we review the most common network structures used in the field and highlight their distinct characteristics (Newman, 2018).

2.1.1 Single-Layer Networks

In this section, we present the common network structures, highlighting their key properties and the typical models used to generate them. These are single-layer networks, which can be used as building blocks for constructing multi-layer networks discussed in the following section.

2.1.1.1 Erdős–Rényi Random Networks

Random networks refer to networks in which connections between nodes are formed randomly. The Erdős–Rényi (ER) model is the typical random network model (Erdős and Rényi, 1959). Each possible edge between a pair of nodes is created at random with a probability p, resulting in a network with a degree distribution that follows a binomial distribution. Random networks are characterized by a low average shortest path length and low clustering coefficients (Watts and Strogatz, 1998; Newman, 2018).

Although they represent a baseline in network science, they lack key properties often observed in real-world networks, such as high clustering (Watts and Strogatz, 1998) and power-law degree distribution (Barabási and Albert, 1999).

2.1.1.2 Scale-Free Networks

Scale-free networks are networks that exhibit a power-law degree distribution, defined as $P(k) \sim k^{-\alpha}$ where α in many real-world networks satisfies $2 < \alpha < 3$ (Boccaletti et al., 2006; Pastor-Satorras et al., 2015). This structural pattern has been widely observed in real-world social networks (Barabási and Albert, 1999; Mislove et al., 2007; Sadikov and Martinez, 2009), as well as in other complex systems such as the World Wide Web and citation networks (Newman, 2018). Scale-free networks are inherently heterogeneous, consisting of a few highly connected nodes (*hub nodes*) while the majority of nodes have relatively few connections (Bianconi, 2018a). The Barabasi-Albert (BA) network is a widely used model for generating scale-free networks (Barabási and Albert, 1999). This model is based on the principle of *preferential attachment* growth, where new nodes are more likely to connect to existing nodes with higher degrees. Specifically, starting with an initial graph, new nodes are added iteratively and connected to existing nodes with a probability proportional to their degrees.

2.1.1.3 Small-World Networks

Small-world networks are defined by the small-world phenomenon, which posits that any two individuals in the network are likely to be connected via a short sequence of intermediate acquaintances (Kleinberg, 2000). These networks exhibit two main properties: a high clustering coefficient and a small average shortest path length (Boccaletti et al., 2006; Telesford et al., 2011). The high clustering feature in small-world networks implies that individuals tend to have strong connections with many mutual contacts, whereas the notion of small shortest path length implies that a spreading entity, e.g., information or disease, can propagate rapidly through the network. The Watts–Strogatz small-world network is a widely used model for generating such networks, capturing the small-world features (Watts and Strogatz, 1998). This model starts with a regular lattice, and then each edge (i,j) is rewired with probability p, known as rewiring probability, to (i,w) where w is randomly chosen such that $w \neq i$, and w is not an existing neighbor. The rewiring probability governs both the clustering and average path length properties. Specifically, a low rewiring probability yields a highly clustered network with longer average path lengths, whereas a high rewiring probability yields more shortcuts, leading to shorter average path lengths but a less clustered network.

2.1.2 Multi-Layer Networks

As mentioned earlier, multilayer networks provide a framework for modeling different types of interactions, each with its own structural properties and dynamics, reflecting more accurate real-world settings. They comprise several interconnected single networks that interact and co-evolve with each other, collectively driving the overall system dynamics. Formally, a multilayer network, denoted as $\mathcal{M}$, consists of L networks, where each network corresponds to a layer $l \in \{1, \ldots, L\}$ and is represented as $G_l = (V_l, E_l)$ with a set of nodes V_l and set of edges E_l. Specifically, a multilayer network can be defined as $\mathcal{M} = (\mathcal{G}, \mathcal{E})$, where $\mathcal{G} = \{G_l | l \in \{1, \ldots, L\}\}$ is the set of networks (layers) and $\mathcal{E}$ denotes the interlayer edges that connect nodes across different layers (Boccaletti et al., 2014).

Multilayer networks have attracted significant attention from researchers across various fields. Studies range from exploring a single diffusion process across multiple layers, such as controlling spreading processes governed by SIS diffusion model (Bernal Jaquez et al., 2020) and epidemic spreading dynamics (Jiang and Zhou, 2018), to coupled diffusion processes, such as the role of social support in mitigating epidemics (Chen et al., 2018) and the interplay of information, opinions, and epidemic diffusion (Dai et al., 2019; Wang et al., 2020; Huang et al., 2021; Yin et al., 2022; Chen et al., 2022b; Fügenschuh and Fu, 2023; Xu et al., 2024). They also have a wide range of applications beyond social systems, such as financial, transportation, and others (Bianconi, 2018b). The interplay and control of these multiple structures are inherently complex and require a deep understanding of their behavior and dependencies.

Multi-layer networks can take several forms. In multiplex networks the layers have the same set of nodes but describe different types of relationships (Bianconi, 2018b; Aleta and Moreno, 2019). For example, in social networks, individuals maintain different types of relationships such as friendships, collaborations, or family ties, which can be captured through multiplex networks. Likewise, they can capture communication over different channels such as email, phone calls, and social media. In the context of information–disease models, multi-layer networks are used to distinguish between the space in which information is transmitted and the space in which disease spreads (e.g., (Dai et al., 2019; Yin et al., 2022)). Multiplex networks are particularly common in this domain, where one layer may represent, for example, a social media platform through which individuals receive health-related information, while another layer captures the physical contact network through which a disease propagates. A detailed discussion is provided in section 2.4.4.

The use of multi-layer networks also extends to time-dependent connections in temporal networks, where each layer represents interactions occurring at a given time, known as multi-slice networks (Bianconi, 2018b). Another class of multilayer networks involves layers with different sets of nodes, known as networks of networks

or interdependent networks. In these networks, each layer has its own set of nodes connected through inter-layer links, often related to dependencies between different systems. One important application of such networks involves scenarios in which nodes in one layer depend on control nodes in another (Boccaletti et al., 2014). An example is infrastructure systems, including power grids, water supply networks, airline systems, and others (Bianconi, 2018b). One example problem is that of cascading failures, where a disruption in one layer impacts the functionality of another. For instance, in the case of a power grid and a communication network, the shutdown of a power station results in a failure of linked communication nodes that rely on electricity (Aleta and Moreno, 2019). Although these are the most common types of multi-layer networks, (Boccaletti et al., 2014) introduces other forms.

In our study, we focus particularly on multiplex networks to model the coupled dynamics of information and disease to capture different types of social contacts for individuals, specifically virtual contacts where information and opinions are exchanged, and physical networks where disease transmission occurs. One significant gap in the literature is the lack of understanding of designing network-based interventions in multilayer settings. In coupled information-disease systems, a key question arises: which layer should be targeted? It is unclear which intervention is more effective, intervening in the information layer or in the disease layer, in containing the spread of disease. For example, if a campaign neglects the virtual communication layer and focuses solely on face-to-face interactions, or vice versa, will it still be effective? These questions highlight the need to investigate how interventions across different layers interact and contribute to the overall system dynamics.

We systematically address the problem by beginning our analysis with a single-layer network to gain insights into controlling the interaction dynamics of information and disease. This allows us to first design and evaluate the effectiveness of various intervention strategies in a simplified setting. We then extend our investigation to examine the influence of the interacting dynamics of distinct network structures in multilayer networks.

2.2 Modelling Opinion Diffusion

Social influence is a process in which individuals are affected by others' actions and tend to behave in a similar way (Guille et al., 2013). Opinions, information, ideas, or behaviors can spread through social networks as a form of social contagion, influencing public behavior. In the context of epidemic control, the propagation of opinions or information can significantly affect the trajectory of an outbreak. For example, the spread of anti-vaccine opinions can lead to a larger epidemic size, as

individuals who adopt such views are less likely to get vaccinated. This not only affects individuals' own protection but also leads to the spread of such attitudes through social interactions, impacting vaccine uptake and hence the overall population immunity. Therefore, understanding the key elements of opinion diffusion and how to model them is essential for capturing and managing the flow of vaccination opinions and dynamics within a social network.

Opinion diffusion models are frameworks that model the propagation of information and opinion adoption within social networks. By integrating these models with network models, we can simulate social systems considering the patterns of social contacts and the propagation of social influence through opinion dynamics. The subsequent sections provide the fundamental principles of modeling the dissemination of opinions across a social network. We first review existing models of opinion diffusion, highlighting their main categories and distinguishing features. We then explore how these models have been applied in relevant studies that addressed coupled opinion–disease dynamics.

2.2.1 Opinion Diffusion Models

Social contagion can spread in a variety of ways within a social network. Generally, social contagion is classified into two types: *simple contagion* and *complex contagion* (Centola et al., 2007; Centola and Macy, 2007; Centola, 2010; Romero et al., 2011). In *simple contagion*, a single exposure is sufficient for an agent to adopt an idea or be influenced by others. In contrast, in *complex contagion*, an agent requires multiple exposures to adopt an idea or be influenced by others. A variety of models have been proposed to describe opinion diffusion processes. Although peer influence is a key component in these opinion diffusion models, their dynamical properties vary significantly. This variability is not surprising, as social influence and the ways individuals interact and form opinions are highly complex. Hence, each model incorporates different assumptions either to simplify the system for analytical tractability or to enhance realism.

Generally, opinion diffusion models are classified into discrete and continuous. Discrete models involve a finite set of opinions that individuals can hold, while in continuous models opinions lie on a continuous scale (Castellano et al., 2009; Dong et al., 2018). Examples of discrete opinion models include the linear threshold (LT) model, the independent cascade model (IC), the voter model, the majority rule model, and the Sznajd model (Kempe et al., 2003; Sîrbu et al., 2017; Castellano et al., 2009). In contrast, continuous opinion models examples include the Bounded confidence model, the Deffuant model, the DeGroot model (Castellano et al., 2009; Dong et al., 2018). In addition to these, epidemiological models used to simulate disease spread, such as the SIR (Susceptible-Infectious-Recovered) and the SIS

(Susceptible-Infectious-Susceptible), have also been utilized to simulate opinion dynamics (Chen et al., 2013).

In the present study, we are interested in discrete opinions, as we investigate the dynamics of vaccination decisions, where individuals choose either to vaccinate or not. The aforementioned discrete models are widely used and well-established in the opinion diffusion field. These models differ in their structural assumptions and can be broadly categorized into progressive and non-progressive diffusion models (Chen et al., 2013). In progressive models, opinions remain unchanged once adopted. In contrast, non-progressive models allow for individuals' opinions to evolve over time and shift back and forth due to their interactions.

Progressive models, in which opinions cannot revert once adopted, are commonly used in scenarios such as modeling innovation technologies adoption (e.g., smartphones), or vaccination acceptance particularly for non-seasonal diseases such as parental decisions regarding the measles vaccine. In such scenarios, once an individual adopts an opinion, it typically leads to a corresponding action such as purchasing a product or, in the case of vaccination, parents vaccinating their children. For measles-like disease, where vaccination provides full immunity (CDC, 2025), progressive models provide a suitable framework for modelling such dynamics (e.g. (Campbell and Salathé, 2013)). However, for diseases requiring repeated vaccination such as flu-like diseases, factors like vaccine efficacy can influence decisions to vaccinate again, making non-progressive models a more appropriate framework as discussed below. The most widely used models in this class include the Linear Threshold Model (LT) and the Independent Cascade Model (IC) (Kempe et al., 2003). In the IC model, each infected (or activated) node has one chance to transmit its influence to each of its inactive neighbors with a given probability (Chen et al., 2013). These nodes cannot exert further influence in subsequent time steps. The process continues until no further activations occur. Such a model exemplifies a simple contagion process, in which nodes can be influenced by one exposure. In the Linear Threshold (LT) model, each node is assigned a random threshold between 0 and 1, and each edge associated with an influence weight (Chen et al., 2013). At each time step, if the sum of the weights for a given node from its active neighbors meets or exceeds its threshold, it becomes influenced and active. This dynamic reflects a complex contagion process, where multiple exposures are required, in which higher threshold values indicate that more active neighbors are needed for activation and vice versa. Recent studies have proposed extensions of these classical models, such as variants of the Linear Threshold (LT) model (Pham et al., 2020; Yang et al., 2019a; Liu et al., 2016).

In contrast, non-progressive models, in which opinions can change back and forth over time, are commonly used in scenarios such as modeling fluctuations in public opinion on topics like climate change, political elections, or vaccination decisions in the context of seasonal diseases such as flu-like disease—particularly with uncertainty

about vaccine efficacy. The most widely used models in this class are the voter model (Clifford and Sudbury, 1973) and SIS model (which will be described later in epidemics models section 2.3). In the voter model, each node in the graph has two states 0 and 1. In each time step, a node selects one of its neighbors at random and mimics the state that this neighbor had in the previous step (Chen et al., 2013). This clearly captures simple contagion dynamics, where influence can be transmitted through a single exposure.

2.2.2 Opinion Models in Coupled Opinion-Disease Dynamics

In coupled opinion-disease models, researchers have employed a range of techniques to model opinion transmission. For instance, epidemic models have been used in (Mehta and Rosenberg, 2020; Mumtaz et al., 2022; Chen et al., 2022b; Yin et al., 2022), while others have employed opinion diffusion models, such as threshold-based model (Campbell and Salathé, 2013), voter-like model (Salathé and Bonhoeffer, 2008), the m-model (Dai et al., 2019), and a variant of Axelrod's model (Dorso et al., 2017). The (Dorso et al., 2017) model incorporated a homophily feature into the social interactions among individuals, which was defined as the tendency for individuals with similar traits to form social connections and to be more likely to interact with each other (McPherson et al., 2001; Dorso et al., 2017).

The utilization of epidemiological models to describe the spread of opinions provides simple contagion, where an individual can be influenced by a single exposure. In contrast, empirical studies have demonstrated that influence propagation exhibits a complex contagion diffusion pattern (Centola, 2010; Romero et al., 2011), where multiple exposures are required to influence an individual. Consequently, opinion diffusion models are better suited to capture the dynamics of social interactions, as they incorporate more realistic behavioral characteristics that govern the adoption of opinions and decision-making in the real world. In our model, we utilize complex contagion due to its consistency with real-world social dynamics. Specifically, building on the model proposed by Campbell and Salathé (2013), we developed an opinion model that incorporates complex contagion dynamics in addition to the dual propagation of negative and positive opinions, as detailed in Chapter 3.

2.3 Modelling Epidemic Spread

Epidemic models are essential frameworks for capturing the dynamics of infectious diseases and gaining insights into how to mitigate their adverse consequences. They can inform policy decisions on potential interventions, such as lockdowns (Doyle et al., 2024), school closures (Viner et al., 2020), travel restrictions (Chinazzi et al.,

2020), or others. The ability of these models to replicate a variety of scenarios aids policymakers in selecting the most appropriate response strategies. They can also guide the optimal allocation of limited resources, such as vaccines (Matrajt et al., 2021; Federico et al., 2024), medical resources (Vo et al., 2021), or others(Nowzari et al., 2015). Furthermore, these models enable the prediction of outbreaks and the undertaking of timely precautionary measures. For example, the epidemic threshold, a critical value that determines the occurrence of an outbreak, can be obtained from the basic reproduction number, R_0. This number represents the average number of secondary infections generated by a single infected individual (Hethcote, 2000; Keeling and Rohani, 2008; Wang et al., 2016). If $R_0 > 1$, the disease invades the population, whereas if $R_0 < 1$, the disease dies out.

Traditionally, epidemic models represent the dynamics of a disease using compartmental structures, in which the entire population is divided into a number of distinct groups based on their health status (e.g., susceptible, infected, recovered) (Hethcote, 2000; Keeling and Rohani, 2008). Transitions between these compartments are governed by predefined rates that describe disease progression. These dynamics are typically modeled using ordinary differential equations (ODEs), assuming a homogeneous and well-mixed population (Hethcote, 2000; Keeling and Rohani, 2008; Pastor-Satorras et al., 2015). Several epidemic models have been proposed in the literature, each tailored to different types of diseases with different levels of resolution in capturing disease dynamics. Although each model has its own set of parameters that describe various aspects of disease progression, two fundamental parameters are common to all: the transmission rate, β, which quantifies the likelihood of disease spread between susceptible and infected individuals; and the recovery rate, γ, which indicates the rate at which infected individuals recover.

In this section, we provide a review of the fundamental epidemic models used in the field. These are the Susceptible-Infected-Recovered (SIR) model, Susceptible-Infected-Susceptible (SIS) model, and Susceptible–Exposed–Infected–Recovered (SEIR) model (Hethcote, 2000; Keeling and Rohani, 2008; Pastor-Satorras et al., 2015; Wang et al., 2016). These models are well-known and widely used in the field of epidemiological research (e.g., (Campbell and Salathé, 2013; Dorso et al., 2017; Yin et al., 2022)). They capture a wide range of disease dynamics: The SIR model accounts for lasting immunity after infection and thus applies to diseases that confer immunity; The SIS model describes disease that does not grant immunity, such as flu; and the SEIR model is similar to the SIR model but considers a latent incubation period occurring before infection. Building on these fundamental models, several variants have been proposed such as the SIRS model which accounts for temporary immunity in which individuals can become susceptible again (Keeling and Rohani, 2008), and the SIRV model which incorporates a vaccinated compartment (Kar and Jana, 2013).

The assumption of a homogeneous and well-mixed population in these compartmental models implies that every individual is equally likely to come into contact with anyone else (Newman, 2002). However, this does not accurately reflect real-world settings. In practice, individuals interact within limited and structured social circles, leading to heterogeneous contact patterns (Newman, 2002; Pastor-Satorras et al., 2015). To better reflect this, networks are used to model this type of interaction as discussed in previous sections. Accordingly, network-based variants of these traditional compartmental models are employed to simulate the disease states of individual agents (Newman, 2002; Keeling and Eames, 2005; Pastor-Satorras et al., 2015). Specifically, the edges between nodes represent potential routes of disease transmission through which infected nodes can spread the disease to their susceptible neighbors with some probability (Newman, 2002; Keeling and Eames, 2005).

In the following sections, the dynamics of each epidemiological model are described in detail.

2.3.1 Susceptible-Infected-Recovered (SIR) Model

The SIR model is a standard epidemic model that is commonly utilized by researchers in epidemic modeling (Newman, 2002). It was first proposed and described by (Kermack and McKendrick, 1927) and has been extensively studied by researchers in the epidemiological field. A recent review can be found in (Wang et al., 2016).

In an SIR model, the population is divided into three groups: Susceptible *(S)*, Infected *(I)*, and Recovered *(R)*, depending on their health status. The Susceptible group includes individuals who are susceptible to an infection from infected people but are not yet infected. The Infected group contains individuals who are infected and are infectious to others. The Recovered group contains individuals who have recovered and are no longer contagious to others. The transition process of the population between the different groups is captured through a set of ordinary differential equations (ODEs) as follows (Keeling and Rohani, 2008):

$$\begin{aligned} \frac{dS}{dt} &= -\beta SI, \\ \frac{dI}{dt} &= \beta SI - \gamma I, \\ \frac{dR}{dt} &= \gamma I, \end{aligned} \tag{2.6}$$

where S, and I represent the proportion of the population in susceptible and infectious states, respectively. In this model, there are two main transitions: $S \rightarrow I$ and $I \rightarrow R$, representing the transition from the susceptible state to the infected state, and from

the infected state to the recovered state, respectively. These transitions are governed by the infection rate, β, and the recovery rate, γ, respectively, where $0 \leq \beta, \gamma \leq 1$.

2.3.2 Susceptible-Infected-Susceptible (SIS) Model

The SIS model stands for Susceptible-Infected-Susceptible and is considered the simplest epidemic model (Pastor-Satorras et al., 2015). Unlike the SIR model, this model does not include immunity, meaning that a node can become infected again after recovering. The SIS model consists of two states: susceptible *(S)* and infected *(I)*. There are two transitions between these states: $S \rightarrow I$, where a node becomes infected and moves from the susceptible state to the infected state, and $I \rightarrow S$, where a node returns to the susceptible state and can be re-infected, as there is no immunity conferred by infection. Similar to the SIR model, these transitions are governed by the infection and recovery rates. The transition process of the population between the different groups is captured through a set of ordinary differential equations (ODEs) as follows (Keeling and Rohani, 2008):

$$\begin{aligned} \frac{dS}{dt} &= \gamma I - \beta SI, \\ \frac{dI}{dt} &= \beta SI - \gamma I, \end{aligned} \tag{2.7}$$

where S, and I represent the proportion of the population in susceptible and infectious states, respectively.

2.3.3 Susceptible–Exposed–Infected–Recovered (SEIR) Model

The SEIR model stands for Susceptible–Exposed–Infected–Recovered model. It is a variant of the SIR model mentioned above, but includes an additional group for the exposed state (E) that account for incubation period of a disease. This state represents individuals who have been exposed to the infection but they are not infectious (Wang et al., 2016; Keeling and Rohani, 2008). The transitions in this model are $S \rightarrow E$, $E \rightarrow I$, and $I \rightarrow R$. The susceptible individual can become exposed to the infection with a transmission rate, β, and move to state E. Individuals in the exposed state transfer to the infected state with a rate α. Infected individuals can infect others with a transmission rate β. Finally, infected individuals recover and move to the recovered state R at recovery rate γ. Once individuals recover, they are assumed to have developed immunity to the disease and cannot be infected again. The transition process of the population between the different groups is captured through a set of ordinary differential equations (ODEs) as follows (Keeling and Rohani, 2008):

$$\begin{aligned}\frac{dS}{dt} &= -\beta SI,\\ \frac{dE}{dt} &= \beta SI - \alpha E,\\ \frac{dI}{dt} &= \alpha E - \gamma I,\\ \frac{dR}{dt} &= \gamma I,\end{aligned} \tag{2.8}$$

where S, E, and I represent the proportion of the population in susceptible, exposed, and infectious states, respectively, and α represents the incubation rate.

2.4 Coupled Epidemics: Interacting Dynamics of Information and Disease

In many real-world situations, disease transmission is not isolated from the flow of information; rather, the two processes are tightly coupled, forming coupled contagion dynamics. The diffusion of health-related information can significantly influence individuals' behaviors and decision-making processes, and can affect the transmission dynamics of diseases within a population. The overwhelming volume of information, its rapid spread, and their associated impact on the spread of diseases represent a major public health concern. Several real-world examples demonstrate the detrimental impact of misinformation on public health efforts to contain disease spread, including Covid-19 (Lee et al., 2022; Pierri et al., 2022), Ebola (Vinck et al., 2019), and measles (Hotez et al., 2020), which also witnessed resurgence recently in 2023 (WHO, 2024).

Information has a dual influence on public health and disease control, as it can not only lead to negative behaviors through the spread of misinformation but also foster preventive behaviors through the propagation of awareness (Alahdal et al., 2020). The interplay between information diffusion and disease spread has received significant attention in the literature (Campbell and Salathé, 2013; Dai et al., 2019; Yin et al., 2022; Sontag et al., 2022; Mumtaz et al., 2022; Chen et al., 2022b; Funk et al., 2009; da Silva et al., 2019; Kabir et al., 2019; Mehta and Rosenberg, 2020; Fügenschuh and Fu, 2023; DeVerna et al., 2025). One set of studies investigates leveraging information to raise awareness and encourage individuals to adopt precautionary measures to control epidemics (e.g., (Funk et al., 2009; da Silva et al., 2019; Kabir et al., 2019; Ruan et al., 2012; Sontag et al., 2022)). Other studies examine the adverse impact of misinformation and its influence on disease prevalence such as its effect on vaccine decision-making (e.g., (Campbell and Salathé, 2013; Mehta and Rosenberg, 2020; Yin et al., 2022; Chen et al., 2022b; DeVerna et al., 2025)). Some explore the competitive diffusion of conflicting information and its effects on disease transmission (e.g., (Dorso et al., 2017; Dai et al., 2019; Huang et al., 2021; Fügenschuh and Fu, 2023)).

In the following, we provide an overview of the modeling approaches commonly used in the field (Section 2.4.1). As this study focuses on vaccine-related information, we also review existing mechanisms for modeling vaccination behavior (Section 2.4.2), followed by a discussion of the impact of vaccine misinformation on disease dynamics (Section 2.4.3). Finally, we review studies that address the interplay between information diffusion and disease transmission by accounting for different diffusion channels with distinct network structures through the use of multi-layer networks. (Section 2.4.4).

2.4.1 Modelling Approaches

Two primary modeling approaches are used in the existing literature to model the diffusion processes of dynamic systems: equation-based modeling and agent-based (or network-based) modeling. Agent-based modeling provides a microscopic view of system evolution by capturing individual-level interactions and accounting for heterogeneous contact patterns (Hunter et al., 2018), thus reflecting more realistic scenarios. However, it is computationally intensive compared to equation-based modeling. In contrast, the equation-based approach offers a macroscopic perspective by approximating the average population behavior using mathematical frameworks such as ordinary differential equations (ODEs) to represent the system, assuming a well-mixed and homogeneous population (Hunter et al., 2018). Both approaches are widely employed in the literature (e.g., (Campbell and Salathé, 2013; Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Fügenschuh and Fu, 2023) for the agent-based approach and (Mehta and Rosenberg, 2020; Mumtaz et al., 2022; Sontag et al., 2022; Harari and Monteiro, 2022) for the equation-based approach).

Furthermore, in coupled information-disease models, the timing of information propagation and disease spread can either be simultaneous, or information may precede the onset of an epidemic outbreak. This mainly depends on the context of the problem being modeled. The former scenario is particularly relevant in cases where information spread can influence the dynamics of an ongoing disease. For example, the dissemination of disease-related information (e.g., health alerts or risk communication) aimed at raising awareness and encouraging protective behaviors (e.g., (Sontag et al., 2022)). Another example is the spread of information related to newly developed vaccines during ongoing outbreaks whether that is misinformation or promotional content such as the development and rollout of Covid-19 vaccines while the epidemic was still in progress (Lee et al., 2022). Conversely, the latter scenario, in which information dissemination occurs at a different stage than disease spread, is particularly relevant in situations where information might impact early protective measures such as influencing vaccination decisions for childhood diseases like measles (Hotez et al., 2020). In such a scenario, vaccination decisions are made

during early childhood, while disease outbreaks often occur later during school age. Several studies in the existing literature adopted this approach (e.g., (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013; Dorso et al., 2017)).

There is no single model or setting that can fully capture the diversity of real-world scenarios. The choice of assumptions and modeling approaches is context-dependent and influenced by several factors such as the nature of the disease, disease seasonality, the type of information being spread, the timing of interactions, the effectiveness of interventions like vaccination, vaccine efficiency, and other contextual variables.

2.4.2 Vaccination Behavior in Coupled Models: Information-Driven vs. Game-Theoretic

The main focus of our study is to explore how behavioral responses to vaccination influence the dynamics of disease spread. Thus, we are particularly interested in coupled dynamics that address the correlation between vaccination behavior and disease spread. Existing models of vaccination behavior can be categorized into information-driven models and game-theoretic models. In information-driven models, vaccination decisions/protective measures are influenced by exposure to information. These studies consider the interaction dynamics between awareness/misinformation/opinion and disease spread (Salathé and Bonhoeffer, 2008; Funk et al., 2009; Ruan et al., 2012; Campbell and Salathé, 2013; da Silva et al., 2019; Mehta and Rosenberg, 2020). In such models, vaccination decisions are modeled using opinion diffusion frameworks (see Section 2.2), where social influence and external influence (e.g., media, public health campaigns) play a significant role in shaping individual choices. For example, (Salathé and Bonhoeffer, 2008) employs a voter-like model, while (Campbell and Salathé, 2013) utilizes a threshold-based model that considers simple and complex contagion. Similarly, (Dorso et al., 2017) uses a model incorporating homophily, where characteristic similarities between individuals drive communication between agents.

In contrast, vaccination decisions in game-theoretic models are strategic and utility-driven, in which individuals estimate the benefits of vaccination and aim to maximize their payoff through their choices. These studies address the problem as a vaccination game using an evolutionary game-theoretic approach (Fu et al., 2011; Cardillo et al., 2013; Zhang et al., 2014; Ichinose and Kurisaku, 2017; Zhang et al., 2017). In the vaccination game, there are generally, but not restricted to, three strategies: get vaccinated, not get vaccinated and get infected, and not get vaccinated and not get infected (Fu et al., 2011; Ichinose and Kurisaku, 2017). The payoff an individual receives depends on the chosen strategy, with each strategy associated with a specific payoff. Choosing not to vaccinate and getting infected by the disease represents the worst strategy, whereas choosing not to vaccinate and avoiding

infection is the ideal strategy. The latter strategy is the main cause of the so-called free rider phenomenon, in which individuals benefit from herd immunity without incurring any risk. This phenomenon has received significant attention in the game-theoretic literature (Fu et al., 2011; Zhang et al., 2014; Ichinose and Kurisaku, 2017; Zhang et al., 2017; Meng et al., 2022).

Some research explores a combination of these approaches (Kabir et al., 2019; Kabir and Tanimoto, 2019; Yin et al., 2022). For example, Kabir et al. (2019) proposed a model that integrates a vaccination game with awareness diffusion. In their study, in addition to vaccination, aware individuals adopt extra self-protective measures such as wearing masks, thereby reducing their likelihood of infection. Similarly, Yin et al. (2022) developed a model where vaccine decisions are modeled as a vaccination game and can also be influenced by negative vaccine-related information.

As this thesis addresses the problem of vaccine hesitancy, which is fueled by the spread of vaccine misinformation as a critical driver, we focus on information-driven studies that explore opinion diffusion and the spread of information as social contagion influencing vaccination decisions.

2.4.3 Impact of Vaccine Misinformation on Epidemic Dynamics

With particular interest in vaccine misinformation, the diffusion of such information can undermine individuals' willingness to get vaccinated, leading to serious public health consequences and influencing the transmission dynamics of infectious diseases. This problem has attracted increasing research attention. Studies have addressed how anti-vaccine information diffusion on social media fuels vaccine hesitancy (Puri et al., 2020), how information diffusion can lead to the formation of polarized clusters, pro-vaccine and anti-vaccine communities, that reinforce internal beliefs and create echo chambers (Schmidt et al., 2018; Mønsted and Lehmann, 2022), making individuals more resistant to changing their beliefs. When online opinion-based clusters result in a corresponding level of vaccination clustering, the likelihood of local outbreaks increases (Salathé and Khandelwal, 2011). Others have examined other perspectives such as how we can leverage social media data to identify online vaccine clusters and its influence on vaccination (Bello-Orgaz et al., 2017), exploring the correlation between online public vaccine attitudes and vaccine uptake (Cheng et al., 2023), and understanding the communication patterns between anti-vaccine and pro-vaccine groups (Yuan et al., 2019).

Numerous studies have explored how the spread of vaccine misinformation shapes vaccine-related opinions and their relationship to the occurrence of outbreaks (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013; Dorso et al., 2017; Dai et al., 2019; Yin et al., 2022; Mumtaz et al., 2022; Chen et al., 2022b; Mehta and Rosenberg, 2020;

Harari and Monteiro, 2022; Fügenschuh and Fu, 2023). These opinions can spread among individuals as a social contagion, significantly affecting vaccination coverage distribution and, consequently, the transmission dynamics of diseases. In this regard, some studies have focused on the dynamics of anti-vaccine opinions while neglecting the influence of pro-vaccine propagation (Campbell and Salathé, 2013; Yin et al., 2022; Mumtaz et al., 2022; Chen et al., 2022b). Others have examined both anti-vaccine and pro-vaccine dynamics, with anti-vaccine dynamics being treated as a contagion process (Mehta and Rosenberg, 2020; Dorso et al., 2017), while other studies considered the propagation of both anti- and pro-vaccine behaviour (Salathé and Bonhoeffer, 2008; Dai et al., 2019; Harari and Monteiro, 2022; Fügenschuh and Fu, 2023).

In more detail, researchers using compartmental dynamic models (i.e., SIR-like models) that integrate the spread of vaccine misinformation with infectious disease transmission have demonstrated that these processes are interlinked, with their interaction significantly altering outbreak dynamics and contributing to disease resurgence (Mehta and Rosenberg, 2020; Mumtaz et al., 2022; Harari and Monteiro, 2022). Specifically, Mumtaz et al. (2022) explored the impact of vaccine misinformation on disease attack rates. They integrated the two contagion processes, each with its own dynamics and reproduction number, which represents the speed of contagion. By varying the reproduction numbers for both processes, they evaluated epidemic incidence and found that even a low reproduction number for misinformation contagion can significantly reduce vaccine confidence and uptake, ultimately leading to a higher attack rate. On the other hand, Mehta and Rosenberg (2020) proposed a model assuming that only anti-vaccine sentiments are transmissible while pro-vaccine opinions are not. They demonstrated that even after long periods without disease outbreaks, the transient dynamics of anti-vaccine sentiment can lead to repeated and easily triggered epidemics. The authors in (Harari and Monteiro, 2022) considered the co-diffusion of pro-vaccine and anti-vaccine sentiments and highlighted the importance of combating misinformation alongside achieving and maintaining high vaccination rates. Their findings showed that tackling misinformation alone, while vaccination rates are low, is insufficient to contain disease dynamics; however, it remains important, as it can reduce the immunization threshold required for disease eradication. These studies employ an equation-based modeling approach and thus primarily focus on understanding the average behavior of the system during the co-diffusion of misinformation and disease, while failing to capture the complexity of social interactions and the heterogeneity in contact patterns.

In contrast, scholars employing a network-based approach gained deeper insights into how social influence affects vaccine coverage within the population, particularly the structure of unvaccinated communities. One significant challenge associated with the spread of anti-vaccine opinions is the emergence of anti-vaccine clusters, groups of

connected agents who refuse vaccination (anti-vaccine agents), (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013; Dorso et al., 2017; Fügenschuh and Fu, 2023). Despite employing different models of vaccine-related opinion diffusion, these studies collectively highlight the role of social influence propagation in driving the emergence of anti-vaccine clusters.

Specifically, Campbell and Salathé (2013) explored how different types of social contagion, namely simple and complex contagion, of anti-vaccine opinions influence the disease transmission and the final epidemic size. They investigated the formation of anti-vaccine communities under both contagion types and found that complex contagion leads to larger clusters of susceptible anti-vaccine individuals compared to simple contagion, ultimately resulting in a greater epidemic size. Similarly, Salathé and Bonhoeffer (2008) addressed the problem considering the dissemination of both pro- and anti-vaccine opinions. They demonstrated how the adoption of anti-vaccine opinions leads to clusters of susceptible individuals, which in turn result in high outbreak levels. Despite high vaccine coverage in the rest of the population, these susceptibility clusters undermine the effectiveness of herd immunity. This was also demonstrated by Dorso et al. (2017), who developed a coupled model for the transmission of vaccination-related opinions and epidemic spread. Their objective was to examine the impact of a small anti-vaccine movement on an existing pro-vaccination population. Moreover, Fügenschuh and Fu (2023) shows how different placement strategies for pro-vaccine agents can counter the spread of anti-vaccine sentiments and thereby help contain epidemic outbreaks.

One main gap in the existing literature is the limited studies that address the impact of combating vaccine misinformation propagation on the dynamics of infectious disease. Therefore, in this thesis, we address the problem of mitigating the diffusion of vaccine hesitancy, embodied as anti-vaccine opinions, through strategic interventions that disseminate pro-vaccine information. More details are provided in section 2.5.2.1.

2.4.4 Coupled Dynamics in Multi-Layer Networks

The different venues of contagion processes, namely information and disease, and the varying structures of networks have recently attracted researchers' attention. In particular, studies have explored how the spread of information through different social platforms, such as social media, can influence the physical transmission of disease in face-to-face networks through multi-layer network frameworks (Dai et al., 2019; Yin et al., 2022; Huang et al., 2021; Chen et al., 2022b; Fügenschuh and Fu, 2023). In such frameworks, each layer captures a different type of relationship between agents. Typically, but not restricted to, one layer represents information diffusion, while another represents physical contact through which disease spreads.

These models have been employed to study different perspectives, including the impact of awareness in encouraging protective measures (Dai et al., 2019), its role in suppressing disease-related rumors (Huang et al., 2021), the impact of vaccine misinformation in vaccination decisions (Chen et al., 2022b; Yin et al., 2022), and the effectiveness of information intervention campaigns in controlling epidemics (Fügenschuh and Fu, 2023). Specifically, authors in Dai et al. (2019) investigate the co-evolution and interaction of disease awareness transmission and vaccine opinions on epidemic dynamics. In a three-layer network model capturing awareness diffusion, vaccine-related opinions, and disease spread, they demonstrated that awareness not only reduces the transmission rate of the disease through the adoption of protective measures but also increases the vaccination rate. Similarly, (Huang et al., 2021) showed the role of awareness spread in managing rumors and controlling epidemic transmission in a two-layer network.

The influence of negative vaccine-related information on epidemic size has been explored in (Yin et al., 2022; Chen et al., 2022b). In particular, Chen et al. (2022b) proposed a two-layer network model to capture the co-diffusion of vaccine misinformation and disease spread. They revealed that when the negative information has a strong influence, it can reduce vaccination coverage and lower the epidemic outbreak threshold, resulting in an increase in the final epidemic size. In contrast, it has a very slight impact on epidemic spreading when the influence of this information is weak. Similar results are demonstrated in (Yin et al., 2022), where authors considered different network structures and incorporated additional factors into the vaccination decision model, including vaccination costs and herd mentality —in which individuals' decisions are influenced by the dominant behavior of the surrounding majority. In a three-layer networks capturing negative vaccine-related information, vaccination behavior, and disease spread, they found that the influence strength of negative information has a more pronounced effect in random networks compared to scale-free networks. Specifically, even when the transmission rate of negative information is high, a low influence strength has only a slight impact on vaccination decisions and epidemic transmission. High vaccination coverage may still be achieved influenced by other factors such as vaccine efficacy and vaccination costs, resulting in a smaller epidemic size particularly in random networks.

Distinct from other studies, Fügenschuh and Fu (2023) explored how different targeted campaigns to mitigate vaccine misinformation on the information layer can curb epidemic outbreaks on the contact layer. They demonstrated that in a scale-free network structure for information diffusion, targeting hub nodes is an effective strategy to combat the spread of anti-vaccine opinions, resulting in a smaller epidemic size. However, although they distinguish between physical and virtual contacts, the model considers disease transmission only through local interactions using a square

lattice, failing to capture the full complexity and heterogeneity of contact networks in real-world scenarios.

The topology of a network whether representing behavioral dynamics or disease transmission can significantly affect how these spreading processes interact in a coupled system. For instance, in correlated layers representing virtual and physical relationships, awareness might reach a substantial portion of individuals when the information layer is denser than the physical layer, i.e., when there are more links in the information layer (Huang et al., 2021) [1]. Similarly, (Yin et al., 2022) demonstrated how the influence strength of information is shaped by the underlying network structure of influence diffusion. Therefore, considering the different network structures underlying the various spreading processes is crucial for understanding the interplay between them.

2.5 Control of Information Diffusion and Epidemics

In this section, we provide an overview of existing work on controlling information and epidemic spread. These two processes have been studied extensively in the literature as single processes, each with its own techniques and strategies for controlling their dynamics. In the field of information control, the well-known areas of influence maximization and influence minimization focus on maximizing the spread of particular information or behavior, or minimizing the spread of undesirable ones. Similarly, many studies have addressed epidemic control through different approaches such as treatment, vaccination, lockdown, information, and others. In the following, we first discuss information control in section 2.5.1, followed by epidemic control in section 2.5.2 including studies that address epidemic control in coupled dynamics in section 2.5.2.1.

2.5.1 Information Control

Information propagation in social networks has a significant impact on shaping public opinion. This is reflected, for example, in the detrimental impact of misinformation on public health efforts (Vinck et al., 2019; Hotez et al., 2020; Lee et al., 2022). Similarly, several empirical studies have investigated the correlation between online attitudes and offline behaviors. For instance, studies have demonstrated the negative impact of online misinformation on vaccination rates (Pierri et al., 2022), the positive correlation between online pro-vaccine attitudes and higher vaccine uptake (Cheng et al., 2023), the impact of online interactions on offline physical activity (Althoff et al., 2017), and

[1] Correlated layers refer to network layers with overlapping connections between virtual and physical networks.

social media's influence on consumer behavior (Nolcheska, 2017). With the advent of social media, information can spread more rapidly and reach wider audiences. Additionally, most recommendation systems on social media platforms are driven by engagement-based algorithms, which can further amplify harmful content, including misinformation and low-credibility information (Narayanan, 2023; Corsi, 2024; Shin and Shin, 2025).

Controlling the spread of information refers to the strategic engagement with the information flow in a network in order to manage its diffusion. Generally, there are two main objectives with opposing goals. The first is to amplify the reach and influence of beneficial information such as promoting public health awareness or educational content. The second is to curb the spread of harmful content such as rumors or misinformation. These objectives are known as the influence maximization and influence minimization problems which will be discussed in the following sections. Controlling information and influence propagation has many practical applications across different domains, such as raising public health awareness (Yadav et al., 2016), combating misinformation (Mena, 2020; Vivion et al., 2022), commercial applications like viral marketing (Razali et al., 2023), economic development (Alshamsi et al., 2018), and more. Below, we present an overview of influence maximization (Section 2.5.1.1), influence minimization (Section 2.5.1.2), and methods for controlling the spread of misinformation (Section 2.5.1.3).

2.5.1.1 Influence Maximization

The influence maximization problem (Kempe et al., 2003) is an intervention approach in which an external controller strategically targets agents and allocates campaign resources to maximize the final number of individuals adopting a specific opinion subject to budget constraints. This problem involves identifying an optimal subset of agents in the network to seed an influence, with the ultimate goal of maximizing its spread. It has attracted significant attention in the literature (Masucci and Silva, 2014; Kandhway and Kuri, 2014; Masuda, 2015; Eshghi et al., 2018; Erkol et al., 2019; Romero Moreno et al., 2021, 2020; Cai et al., 2021), with various targeting methods proposed considering different diffusion models. For instance, Kempe et al. (2003); Erkol et al. (2019) investigate the problem within the independent-cascade model, Romero Moreno et al. (2021); Cai et al. (2021) employ the voter model, and (Kandhway and Kuri, 2014) explores the problem using epidemic models.

Existing studies on influence maximization have addressed several aspects of the problem. For example, influence maximization with regard to time constraints has been explored in (Cai et al., 2021), investigating how to maximize the spread of influence within a specific time horizon. Other works have considered the competitive influence problem where more than one external controller aims to spread its

influence in the network (Masucci and Silva, 2014; Chakraborty et al., 2019; Romero Moreno et al., 2021). Furthermore, the effect of noise and uncertainty in the diffusion process has been investigated in (Brede et al., 2019), highlighting more realistic settings. In addition, a systematic comparison aimed at quantifying the performance of existing influence maximization methods across different real-world networks was conducted in (Erkol et al., 2019).

Scholars in this area have proposed a variety of methods to address the problem of whom to target in a network and how to optimally allocate control resources to maximize the spread of influence. Heuristic methods are one of the common approaches in which the selection of nodes depends either on the network structure, such as high degree, betweenness centrality, or other network metrics (Erkol et al., 2019; Alshamsi et al., 2018), or on local network information, such as opponents' targeting strategies (Romero Moreno et al., 2021). Optimization-based methods are another approach that applies optimization techniques to solve the influence maximization problem. In this regard, as solving the problem using optimization is NP-hard (Kempe et al., 2003), greedy algorithms have been widely used to find near-optimal sets of nodes with approximation guarantees (Kempe et al., 2003; Chen et al., 2009). These algorithms forecast the impact of each node in maximizing the influence and select those that are the most influential. In addition, an adaptive targeting strategy using gradient ascent has been proposed in (Romero Moreno et al., 2021), where the controller adjusts its strategy dynamically with the evolution of the diffusion dynamics.

2.5.1.2 Influence Minimization

In contrast to influence maximization, influence minimization aims to block the diffusion of specific information, with the ultimate goal of minimizing the final number of individuals influenced by this information. Its primary objective is to minimize the propagation of undesirable influence within a social network by identifying a subset of agents or edges to immunize. This problem is commonly referred to as misinformation mitigation and has also attracted a significant body of literature focused on combating the spread of misinformation (Budak et al., 2011; Zhang et al., 2015; Liu et al., 2016; Tong et al., 2018; Yang et al., 2019a). The solutions to this problem typically involve either blocking certain nodes (Pham et al., 2020) or edges (Zareie and Sakellariou, 2022), or applying true information campaigns (Budak et al., 2011; Zhang et al., 2015; Liu et al., 2016; Tong et al., 2018; Yang et al., 2019a). For a detailed overview of this topic see (Zareie and Sakellariou, 2021).

Both heuristics and optimization methods have also been employed in the field. For example, network structure attributes have been used to select the k highest score nodes to seed the true information using Pagerank and degree centrality (Yang et al.,

2019a; Budak et al., 2011). Heuristic algorithms based on local network information have also been introduced. For example, Budak et al. (2011) suggested the Early Infectees and Largest Infectees heuristics, in which the former seeks nodes that are susceptible to infection, whereas the latter seeks nodes that would transfer the infection to the greatest number of nodes if they become infected. Additionally, Zhang et al. (2015) proposed a Dominating Influence (DI) heuristic method to first identify the nodes which will spread their influence to the greatest number of nodes, then select the candidate seed sets for the true campaign from nodes that are one step away from those gateway nodes who contribute most to misinformation containment and potentially increase good information spread. In Yang et al. (2019a) a heuristic method has been proposed to compute the contribution of the influence spread for each node adopting a rumor based on the number of its out neighbors, then selecting the first k contributors rated according to their contributions to rumor propagation. Greedy algorithms have also been utilized to find the sub-optimal seed set to seed true information (Tong et al., 2018). In their study, the authors utilize the greedy algorithm to select potential influential nodes and investigate a scenario of a multi-cascading diffusion model that incorporates the concept of cascade priority, which determines individual responses to received cascades.

Combating misinformation diffusion extends beyond network-based strategies to include broader approaches. In the next section, we provide an overview of existing strategies to suppress the diffusion of misinformation.

2.5.1.3 Misinformation Control

Intervention strategies to combat misinformation in social networks represent a broad field that encompasses multiple dimensions and objectives. One dimension involves content-related strategies, which focus on the type of information delivered to the population. In this context, various methods have been proposed in the literature, including prebunking, debunking, fact-checking, and others (Roozenbeek et al., 2020; Chen et al., 2023; Denniss and Lindberg, 2025). Prebunking, for example, is based on inoculation theory and inspired by the concept of medical immunization; it aims to increase individuals' resistance to future misinformation by exposing them to weakened doses of misinformation content in advance (Compton, 2013; Roozenbeek et al., 2020). Debunking, by contrast, is a corrective approach that seeks to replace false beliefs with accurate information (Tay et al., 2022). These interventions have been widely studied and demonstrated to be effective in combating the influence of misinformation (Roozenbeek et al., 2020; Vivion et al., 2022; Roozenbeek et al., 2022; Vraga and Bode, 2018; Seo et al., 2022; Bruns et al., 2024).

Another important dimension of intervention strategy addresses the question of whom to target within a network. These network-based interventions represent a

different class of strategies that aim to strategically identify and target a subset of individuals in the network to seed or block influence diffusion for efficient control which represents influence maximization and influence minimization problems as discussed in sections 2.5.1.1 and 2.5.1.2, respectively. This type of intervention is the primary focus of our study. In this context, as we have seen, scholars have explored and developed targeting methods to solve these two problems based on network structural properties or agents' local information as well as leveraging optimization techniques.

Both content-related and network-based approaches play complementary roles, and their integration has the potential to further enhance the effectiveness of a campaign. Content-related interventions such as prebunking and debunking are beyond the scope of this study. However, both approaches could be integrated with our targeting approach by tailoring the campaign content for the selected individuals.

Leveraging information control techniques, i.e., influence maximization and minimization methods, to strategically mitigate vaccine hesitancy has the potential to manage the spread of vaccine-preventable diseases, as information can positively influence epidemic dynamics and help suppress its spread. Despite this potential, the application of these techniques within coupled models aimed at controlling epidemic spread has received limited attention in the literature.

Moreover, with particular interest in the misinformation mitigation problem, studies in this field have mainly concentrated on reducing the number of individuals who are negatively influenced by misinformation. However, considering the interacting dynamics within coupled information-disease processes, reducing the number of negative opinion adopters alone may be insufficient to suppress an epidemic. The distribution pattern of individuals who adopt negative opinions in the social network as well as their location and connectivity within different clusters or regions, are key factors that can significantly impact the extent of the epidemic. As discussed in section 2.4.3, one of the key challenges of vaccine misinformation propagation is the formation of anti-vaccine communities composed of individuals who refuse vaccination, representing vulnerable areas for disease spread. Therefore, in this thesis, we not only aim to reduce the final number of agents adopting negative opinions but also consider their spatial coalescence in the network.

2.5.2 Epidemic Control

Controlling the transmission of infectious disease and preventing epidemic outbreaks are a central objective of public health administrations worldwide. To this end, both pharmaceutical interventions, such as vaccination and antiviral treatments, and non-pharmaceutical interventions, such as social distancing, quarantine measures,

and public awareness campaigns are employed. However, the effectiveness of these interventions is significantly affected by public perception and behavioral responses (Teasdale et al., 2014; Vinck et al., 2019; Hotez et al., 2020). For example, disease-related rumors or vaccine misinformation can influence people's willingness to adopt protective measures or accept vaccination. Therefore, increasing attention has been given to the role of public behavior in epidemic control, and numerous studies have demonstrated the impact of information dissemination in either mitigating or exacerbating epidemic outbreaks, as discussed in the previous section.

Existing studies have examined the impact of various interventions in controlling epidemic outbreaks (Wang and Lan, 2023; Paireau et al., 2023; Zhang et al., 2023; Li et al., 2022). For example, several studies investigate the effect of different interventions on reducing COVID-19 transmission, including testing (Wang and Lan, 2023), lockdowns, school closures, and curfews (Paireau et al., 2023), as well as mobility restrictions (Li et al., 2022). Similarly, (Zhang et al., 2023) shows that a broad set of interventions, ranging from full lockdowns to mask wearing and temperature monitoring, is associated with significant reductions in incidence across multiple infectious diseases, such as influenza and measles, while also contributing to lower economic costs.

Theoretical research has also explored epidemic control through different interventions, such as treatment, vaccination, and isolation (Kar and Jana, 2013; Mondal and Khajanchi, 2022; Sun et al., 2023; Federico et al., 2024). Other studies have considered different perspectives, such as identifying cost-optimal strategies under budget constraints (Nowzari et al., 2015), developing cost-effective lockdown strategies (Doyle et al., 2024), or minimizing testing costs (Acemoglu et al., 2024). The majority of these studies address the problem by introducing control variables into compartmental epidemic models represented as systems of ordinary differential equations (ODEs). These include the classic SIR model (Acemoglu et al., 2024), the age-stratified SEIR model (Doyle et al., 2024), as well as variants of SIR such as SIRV (Kar and Jana, 2013), SAIQJR (Mondal and Khajanchi, 2022), and SIRS (Federico et al., 2024).

In this work, we are particularly interested in controlling epidemic outbreaks by influencing behavioral responses to vaccination, with a particular focus on mitigating vaccine hesitancy. Therefore, the rest of this section is dedicated to relevant studies that address epidemic control through information-based interventions.

2.5.2.1 Epidemic Control in Coupled Dynamics

Despite significant attention to disease-information coupled dynamics, there have been limited efforts to explore the impact of mitigating the diffusion of vaccine

hesitancy on the spread of infectious diseases (Dorso et al., 2017; Ancona et al., 2022; Fügenschuh and Fu, 2023). Further research is needed to design effective intervention strategies to counter the spread of anti-vaccine social influence within coupled diffusion processes and capture its impact on disease prevalence. As discussed in Section 2.5.1, previous studies on information control have thoroughly investigated control strategies within single diffusion processes for the influence maximization problem (Romero Moreno et al., 2021, 2020; Erkol et al., 2019) and the misinformation mitigation problem (Budak et al., 2011; Zhang et al., 2015; Yang et al., 2019a). However, the design of efficient targeting strategies that account for the complex interplay between information diffusion and disease spread in coupled processes remains limited.

Controlling epidemic spread by managing information dissemination through external campaigns within coupled dynamics has emerged as an area of growing research interest (Dorso et al., 2017; Yin et al., 2023; Fügenschuh and Fu, 2023). For example, (Dorso et al., 2017) proposed an intervention scheme to control disease spread by governing the spread of anti-vaccine opinions, assuming that agents who adopt anti-vaccine views can be converted back to pro-vaccine views through dedicated influence. Similarly, Yin et al. (2023) examined the effects of an intervention campaign aimed at promoting preventive behaviors, specifically by increasing social distancing between agents. A recent study that follows an approach similar to our work Fügenschuh and Fu (2023) investigated the impact of different one-time targeting heuristics for placing pro-vaccine seeds on epidemic size in multi-layer networks. In contrast to our dynamic framework, which considers continuous negative influence and both static and adaptive targeting strategies, they considered a static setting in which pro- and anti-vaccine seeds are selected at the beginning of the diffusion process, as we will discuss in Chapter 3. Additionally, another study Ancona et al. (2022), with a similar goal of mitigating vaccine hesitancy proposed strategic targeting methods as a single diffusion process, though it did not consider the effects of these interventions on the dynamics of disease spread.

One significant challenge associated with vaccine refusal, as mentioned above, is the geographic clustering of unvaccinated individuals. Existing literature highlights this geographic clustering (Dubé et al., 2015), as well as the formation of distinct negative clusters in online social networks (Salathé and Khandelwal, 2011). The clustering of unvaccinated individuals undermines herd immunity and creates vulnerable zones for disease transmission (Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Gromis and Liu, 2022). Despite many studies emphasizing the risk posed by anti-vaccine communities and their correlation with the extent of disease spread, a closer examination of the literature reveals the need for research focused on mitigating the expansion and connectivity of anti-vaccine clusters. Understanding the network

structure and connectivity patterns of these emerging communities is therefore critical for effectively containing disease spread.

2.6 Summary

While the role of information dynamics in shaping disease transmission is well-established in the literature, little is known about the effectiveness of strategic information-based interventions aimed at mitigating the spread of vaccine misinformation in controlling epidemics. This knowledge gap raises important questions: How effectively can targeting strategies designed to control information diffusion perform within the coupled dynamic of disease and information spread? How can we design intervention strategies that effectively combat the diffusion of vaccine misinformation? In addition, it remains unclear whether aiming to reduce the number of individuals influenced by vaccine misinformation is sufficient to suppress epidemic outbreaks, and whether it helps prevent the geographic clustering of unvaccinated individuals.

Although many studies highlight the risk posed by anti-vaccine communities and their correlation with the extent of disease spread, there is a lack of research addressing this problem. To fill these gaps, inspired by information control techniques, this thesis explores how different strategies for limiting the spread of vaccine misinformation contribute to curbing epidemic outbreaks. In addition, it addresses the problem of how to disrupt the connectivity of emerging anti-vaccine communities to prevent the geographic clustering of unvaccinated individuals. It also explores the impact of the underlying network structure on the co-evolution dynamics of information and disease and how this impacts the effectiveness of control strategies.

In addition, although considerable research has studied disease dynamics in multi-layer networks with distinct social contact patterns for different spreading processes, misinformation mitigation strategies in these frameworks remain largely unexplored. Specifically, social interactions can take place in both virtual and physical spaces. This raises the question of which space should be targeted to effectively mitigate misinformation. Additionally, whether targeting virtual networks provides more effective control than face-to-face networks remains unclear. In this regard, this thesis presents a systematic and comprehensive investigation of interventions applied across different types of social networks.

We systematically address the problem by beginning our analysis with a single-layer network to gain insights into controlling the interaction dynamics of information and disease. This allows us to first evaluate the effectiveness of various intervention strategies in a simplified setting. We then progressively explore the role of network

topology on these outcomes, and finally extend our investigation to the influence of multi-layer structure.

Chapter 3

Heuristic-based Intervention Campaigns

This chapter has been published as a journal paper in PLOS One (Alahmadi et al., 2025b).

3.1 Introduction

As discussed in Section 2.4, considerable research has examined how the dynamics of vaccine-related information and opinion diffusion impact epidemic spread (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013; Dorso et al., 2017; Dai et al., 2019; Yin et al., 2022). However, few studies have explored how targeted intervention campaigns that aim to mitigate the diffusion of anti-vaccine influence affect epidemic dynamics. This gap persists despite rich literature on strategic targeting in other fields such as influence maximization (Romero Moreno et al., 2021, 2020), and misinformation mitigation (Budak et al., 2011; Liu et al., 2016; Tong et al., 2018; Yang et al., 2019a). Therefore, understanding how different intervention methods for mitigating vaccine hesitancy influence epidemic control remains an open question that requires further investigation.

The problem of controlling epidemic spread by managing the dissemination of information and opinions through external interventions within coupled dynamics is an active research area (Dorso et al., 2017; Yin et al., 2023; Fügenschuh and Fu, 2023). For example, Dorso et al. (2017) proposed an intervention scheme to control disease spread by governing the spread of anti-vaccine opinions, assuming that agents who adopt anti-vaccine views can be converted back to pro-vaccine views through dedicated influence. However, their model considers only anti-vaccine opinions as a form of social contagion, overlooking the diffusion of pro-vaccine influence. Similarly,

Yin et al. (2023) examined the effects of an intervention campaign aimed at promoting preventive behaviors by increasing social distancing between agents. They consider the spread of positive behavior through awareness propagation. Closer to our work, Fügenschuh and Fu (2023) investigated the impact of different static targeting heuristics for placing pro-vaccine seeds on epidemic size in a multi-layer setting, considering the competitive diffusion of both anti- and pro-vaccine opinions. They explored several topology-based heuristics, such as targeting hubs, high betweenness nodes, low-degree nodes, and others, to seed pro-vaccine agents and mitigate disease outbreaks. However, they addressed the problem in a static setting, where both anti- and pro-vaccine seeds were selected at the beginning of the diffusion process. In contrast, our study considers a dynamic setting, in which negative information continuously originates from external sources that randomly influence the network throughout the opinion diffusion phase. Additionally, another study Ancona et al. (2022) with a similar goal of mitigating vaccine hesitancy proposed strategic targeting methods as a single diffusion process, though it did not consider the effects of these interventions on the dynamics of disease spread.

Moreover, the problem of mitigating the propagation of misinformation has received researchers' attention in the field of influence minimization. The influence minimization problem focuses on minimizing the propagation of undesirable influence in a social network. Researchers have addressed this problem either by blocking specific nodes (Pham et al., 2020) or edges (Zareie and Sakellariou, 2022) in the network, or by applying true information campaigns (Budak et al., 2011; Liu et al., 2016; Tong et al., 2018; Yang et al., 2019a, 2020; Zhang et al., 2015). However, these strategies address the problem as a single diffusion process. In the context of vaccine misinformation and coupled information–disease dynamics, understanding how mitigating vaccine misinformation propagation affects disease spread dynamics is vital. Additionally, understanding how different intervention strategies influence the formation and emergence of unprotected (anti-vaccine) communities and, consequently, the resulting epidemic size is crucial.

To bridge these gaps, this chapter explores and proposes different heuristics for seeding pro-vaccine influence to mitigate the spread of anti-vaccine opinions and assess their impact on the final epidemic size. In this context, we consider two approaches for the seed set selection: static and dynamic. In the traditional static approach, the seed set is chosen at the beginning of the campaign launch and remains unchanged throughout the campaign (e.g., (Budak et al., 2011; Yang et al., 2019a; Fügenschuh and Fu, 2023)). In contrast, in the dynamic approach, the seed set is selected repeatedly in different rounds based on specific criteria (Wijayanto and Murata, 2019; Shi et al., 2019). We propose a novel dynamic approach where we update the target set with a new target set based on certain criteria. The selection criteria in this approach are based on the local presence of negative vaccine-related

information. This is an adaptive approach that responds to the evolving dynamics of anti-vaccine information propagation.

Specifically, in this chapter, we make the following contributions to address research questions 1a, 1b, and 1c detailed in Section 1.1. First, we develop a computational framework to examine the effects of mitigating anti-vaccine opinions on the size of epidemics. This model integrates pro- and anti-vaccine opinion diffusion, disease spread, and external vaccine information sources. Second, we explore and compare different heuristics based on network topology and vaccine-related information, aimed at identifying a set of nodes within the network to seed positive influence. We examine their effectiveness in counteracting the diffusion of anti-vaccine opinions and, consequently, in containing the spread of disease. Third, we investigate the impact of resource constraints on these network-targeting strategies, highlighting the trade-off between the number of targeted nodes and the positive influence budget across different campaign types.

The chapter is structured as follows: Section 3.2 describes the framework and model, including opinion diffusion and epidemic spread. Section 3.3 introduces the proposed intervention strategies. Section 3.4 presents the corresponding results. Section 3.4.5 explores the impact of the target set size on campaign performance, comparing this parameter across static and dynamic approaches. Section 3.4.6 provides a comprehensive comparison of the campaigns, considering key parameters such as campaign budget and social influence. Finally, Section 3.5 summarizes the main findings.

3.2 Model Description and Methods

In this study, we consider a scenario where opinion exchanges and the vaccination process occur prior to the spread of the disease. This type of modelling is particularly relevant for diseases like childhood illnesses, e.g., measles, where vaccinations typically take place during the early years of childhood, and the disease manifests during school or preschool age. As a relevant real-world example, after almost two decades of significant progress in global vaccination programs, measles resurged in 2019 (Hotez et al., 2020) and again in 2023 (WHO, 2024) due to a significant decline in vaccine coverage attributed to vaccine misinformation. This disease remains an ongoing global threat, with outbreaks continuing to emerge in various regions (Do and Mulholland, 2025).

Motivated by the model presented in Campbell and Salathé (2013), we developed a framework consisting of a two-stage network-based (agent-based) model, where nodes represent individuals, also referred to as agents and used interchangeably with individuals. The first stage involves opinion diffusion and vaccination processes in

relation to individuals' opinions about vaccines, and the second stage is the disease spread among unvaccinated individuals. This modeling approach is a common methodology found in the existing literature on vaccination programs for childhood diseases (Salathé and Bonhoeffer, 2008) and flu-like diseases (Ichinose and Kurisaku, 2017; Zhang et al., 2017; Fu et al., 2011; Tatsukawa et al., 2021). We assume that the vaccination provides full immunity, implying that all who receive it are immune to the disease. For vaccine opinion diffusion, motivated by the approach of Campbell and Salathé (2013), we developed a model for the dual propagation of positive and negative opinions. Similarly, as in Campbell and Salathé (2013), we utilized the SIR model for the disease spread stage, originally developed by Kermack and McKendrick (1927), which is widely used by researchers in epidemic modeling.

The experiment workflow commences with the stage of opinion diffusion, during which sentiments related to vaccines spread in the network, subsequently influencing individuals' vaccination decision-making. As influenced individuals adopt particular opinions, their social influence starts to propagate through the network. This social contagion inevitably leads to the formation of homogeneous communities, i.e. a community with a particular attitude towards vaccination, with anti-vaccine communities being our primary concern. In our model, we utilize complex contagion, which implies that agents need multiple exposures to be influenced, as this phenomenon has also been observed in influence propagation processes (Centola, 2010; Romero et al., 2011). Following the stage of opinion dissemination, the vaccination process takes place. Subsequently, the infectious disease begins to spread among unimmunized individuals. The details of each stage will be explained in the following sections.

We conduct our experiments using the Watts-Strogatz small-world network model (Watts and Strogatz, 1998). Small-world network provides an effective framework for modelling complex systems, as many real-world networks including social networks exhibit the small-world property (Watts and Strogatz, 1998; Kleinberg, 2000; Barrat and Weigt, 2000; Latora and Marchiori, 2001). In addition, as the model is stochastic in nature, we conducted a number of simulations per scenario to obtain the average epidemic size. In each simulation, we generate the contact network, followed by the diffusion of opinions and vaccination, and finally the disease spread. An illustration of the model stages is shown in Fig.3.1.

3.2.1 Opinion Diffusion

Our model consists of a network G composed of N nodes, represented by $G(V, E)$, where V is the set of nodes representing individuals, $V = \{1, 2, ...N\}$, and E is the set of edges representing contacts between individuals (agents). In this chapter, we assume that this contact network is the same for both the flow of information and the

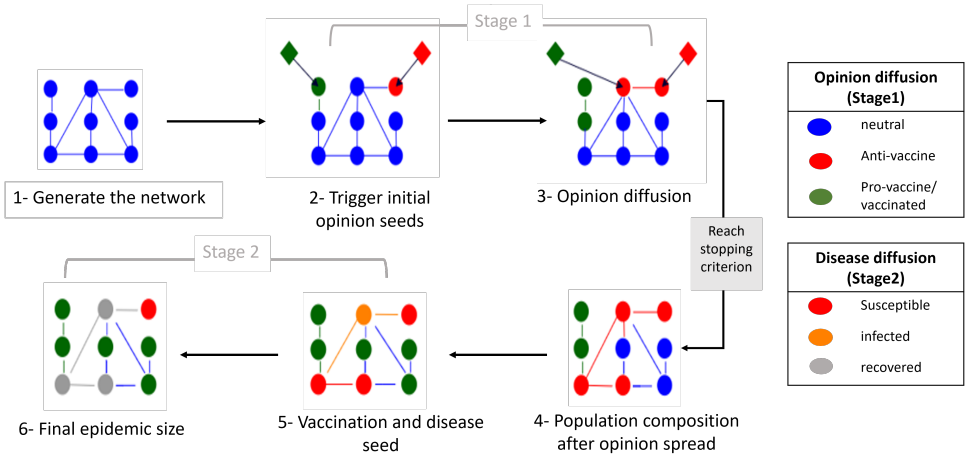

FIGURE 3.1: Illustration of the model describing opinion formation and disease propagation. Blue circles represent neutral individuals, green circles represent pro-vaccine individuals, and red circles represent anti-vaccine individuals. The first stage involves the generation of the social network and the initialization of agent opinion states as agents with neutral opinions. Then, external exposures to positive and negative information triggers the initial seed sets for both anti-vaccine and pro-vaccine contagion. Opinion diffusion continues until a stopping criterion is reached. In this stage, a vaccination takes place for all non-negative individuals. Subsequently, a randomly chosen non-vaccinated individual is infected, and the spread of the disease continues until no further newly infected agents are generated. Finally, we record the number of recovered agents to measure the epidemic size.

transmission of the disease. While this assumption does not fully reflect the complexity of real-world social structures, we consider this simplified structure to gain theoretical insights and understanding of how network topology influences the performance of intervention campaigns. We consider two types of vaccine-related exposures: positive, which spread positive sentiment, and negative, which spread negative sentiment. In addition, we consider exposures from external sources, referred to as general exposure or campaigns, occurring with probabilities μ^- (negative) and μ^+ (positive), as well as through social communication where influence is exerted by opinion adopters on each of their neighbours with probabilities ω^- (negative) and ω^+ (positive), per timestep. Each agent has their own set of counters $\{\phi^-, \phi^+\}$, where ϕ_i^- quantifies exposures to negative, and ϕ_i^+ quantifies exposures to positive sentiments experienced by each agent. Furthermore, θ is an opinion decision threshold that represents an individual's sensitivity to influence. We assume that an agent shifts its opinion from a neutral state to either negative or positive when it has been exposed to θ more exposures of a particular influence.

Each agent i, $i = \{1, 2, .., N\}$, may adopt one of three opinion states $s_i \in \{o^-, o^0, o^+\}$, where o^- is negative, o^+ is positive, and o^0 is neutral. We assume that once an agent changes its state from a neutral to a negative (or positive) state, it remains in that state. Specifically, we assume that individuals who adopt a pro-vaccine opinion choose to immunize, achieving full immunity— a scenario particularly relevant for measles-like

illnesses, where vaccination typically provides full immunity (CDC, 2025). Conversely, anti-vaccine individuals refuse vaccination and maintain their stances. This is consistent with real-world observations of the growing polarization in the vaccine debate, driven by misinformation and social media echo chambers, which reinforce entrenched views (Schmidt et al., 2018; Mønsted and Lehmann, 2022), making opinion changes less likely. In addition, previous studies have shown that correcting false beliefs is challenging once individuals have been exposed to misinformation (Roets et al., 2017). To summarize, following a setup informed by the approach taken in Campbell and Salathé (2013), our model operates in discrete time steps as follows:

- At time $t = 0$, all agents are neutral, i.e., $s_i = o^0, \forall i \in V$.
- At each time step t, negative general exposure exerts influence on the population with probability μ^- per individual. Similarly, positive general exposure exerts influence on the targeted population with probability μ^+ per individual. In addition, each agent i with state $s_i \in \{o^-, o^+\}$, exerts influence on each of its neutral neighbors with probability ω^- for negative opinion adopters and ω^+ for positive opinion adopters. Exposures for an agent i, i.e, $\{i \in V | s_i = o^0\}$, is added to ϕ_i^- if it is negative or added to ϕ_i^+ if it is positive.
- At each time step t, each neutral agent i, i.e, $\{i \in V | s_i = o^0\}$, updates its opinion state as follows:

$$s_i = \begin{cases} o^- & \text{if } \phi_i^- - \phi_i^+ \geq \theta, \\ o^+ & \text{if } \phi_i^- - \phi_i^+ \leq -\theta, \\ o^0 & \text{otherwise.} \end{cases} \tag{3.1}$$

The above process is repeated for τ steps. A visual depiction of external campaigns and social influence is shown in Fig.3.2(A). Next, we make the assumption that all agents with a non-negative opinion will receive the vaccine, while those with a negative opinion will refuse it.

3.2.2 Epidemic Spread

In this phase, disease spread is modeled using the SIR model, with initial infection seeds randomly selected from individuals who adopt an anti-vaccine opinion. Each infected agent can transmit the disease with a probability β per contact per time step, and can recover with a probability γ per time step. The process continues until there are no more infected individuals, and the epidemic size S_r is recorded. S_r is defined as the total number of individuals who experience an infection during the course of the epidemic (House et al., 2013).

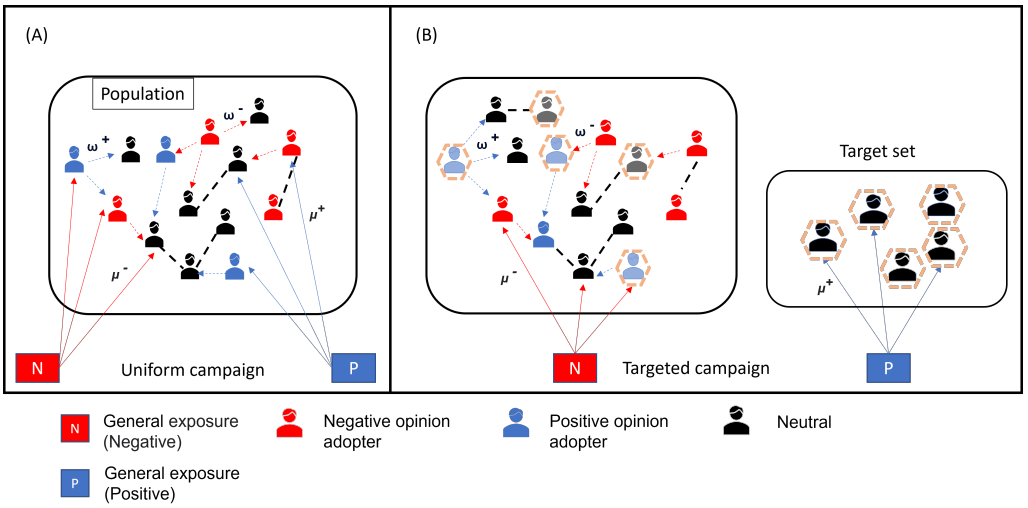

FIGURE 3.2: Illustration of opinion propagation and campaigning methods. The figure shows the exchange of vaccine-related opinions and external exposures, as well as the positive campaign types. (A) Random dissemination of negative and positive vaccine-related sentiments from external campaigns to the public. (B) Targeted positive campaign. μ^-, and μ^+ are the general exposure rates for negative and positive sentiments, respectively. ω^-, ω^+ are the social exposure rates for negative and positive opinions, respectively.

Table 3.1 below presents a list of the model parameters and their corresponding descriptions.

TABLE 3.1: Model parameters and descriptions.

Parameter	**Parameter Description**
N	Population size
μ^-	Negative general exposure rate
μ^+	Positive general exposure rate
ω^-	Negative social rate
ω^+	Positive social rate
ϕ^-	Negative exposure counter
ϕ^+	Positive exposure counter
θ	Opinion formation threshold
T	Target set size
t_r	Update time interval for the dynamic control
τ	Opinion diffusion time steps (stopping criterion)
ζ	Target number of anti-vaccine neighbours
Z	Target number of neutral neighbours
β	Disease infection rate
γ	Disease recovery rate
I_0	Initial number of infected individuals
S_r	Epidemic size

3.3 Campaign Strategies

In this study, we define a positive campaign as a strategic allocation of the strength of positive external vaccine information, denoted as μ_i^+, to the agents $i = 1, ..., N$, with $1/N \sum_i^N \mu_i^+ = \mu^+$. We compare the effectiveness of various types of such positive campaigns against a random negative campaign that spreads negative vaccine information, assuming that each agent can be negatively influenced with an influence strength $\mu_i^- = \mu^-$ at each time step. This section outlines the proposed selection strategies for the target set for the positive campaign. Table 3.2 below provides an overview of the proposed strategies.

TABLE 3.2: Campaign strategies and descriptions.

Campaign type	**Campaign name**	**Acronym**	**Description**
Baseline	Random	StatRandAll	All agents can be influenced with a uniform positive allocation μ^+ at each time step.
Static	Targeted random	StatRandT	An unchanged random subset of the entire population.
	Centrality-based	StatCentT	The T most central agents in the network.
Dynamic	Dynamic random	DynRandT	A random subset of the entire population, which is replaced with a new target set at random every t_r.
	Local information based	DynAntiT	Agents who have at least one anti-vaccine neighbor.
	Advanced LocInfo with single-objective	DynLocT	The T agents with the lowest score according to the number of adjacent anti-vaccine neighbors, Eq 3.2.
	Advanced local-info with multi-objective	DynAdvLocT	The T agents with the lowest score according to the number of adjacent anti-vaccine and neutral neighbors, Eq 3.3.

3.3.1 Random Campaign

In this scheme, we extend the work introduced by Campbell and Salathé (2013), in which they explored the impact of only anti-vaccine opinions on the formation of anti-vaccine communities. In their model, all individuals can be influenced by a general negative exposure with a specific rate. In our baseline campaign, all individuals are also exposed to a pro-vaccine exposure (referred to as positive exposure), with a positive allocation $\mu_i^+ = \mu^+$ at each time step. The campaign scheme is illustrated in Fig.3.2(A).

3.3.2 Targeted Campaigns

In this approach, we aim to target a certain set of agents with neutral opinions, as illustrated in Fig.3.2(B), to efficiently mitigate the spread of anti-vaccine influence. This puts a greater emphasis on specific individuals and has been demonstrated as an effective method, as evidenced by empirical studies (Konstantinou et al., 2021). There are several ways to target them: (i) random selection as a common and intuitive approach; (ii) based on their topological position on the network; (iii) based on their neighbourhood status with regards to local information about vaccine opinions. Each of these will be explained in detail in the subsequent sections.

Furthermore, this approach involves two types of campaigning: static and dynamic. In the static approach, the target set is selected based on predetermined criteria prior to the launch of the campaign and this set remains unchanged. In the dynamic approach, the initial targets are selected at random, however, every t_r opinion updates the target set is updated and replaced by new targets.

In the targeted approach, let T_i =1 if agent i is targeted and $T_i = 0$ otherwise. Accordingly, the positive campaign allocation $\mu_i^+ = T_i \times \mu^+ \times N / \sum_i^N T_i$. This campaigning scheme directs the positive influence budget toward specific individuals rather than distributing it across the entire population, optimizing resource use and increasing the allocation dedicated to targeted individuals. In the following, we have explored and compared the following six heuristics for selecting the target set.

3.3.2.1 Static Campaigns

The following strategies outline the criteria for selecting the target set in the static campaign approach.

1. **Targeted random strategy (StatRandT)**: a random subset of the entire population is selected as the campaign's targets.
2. **Targeted centrality-based strategy (StatCentT)**: this is a topology-based campaign in which targets are selected based on betweenness centrality. The T most central agents in the network are targeted. In the case of ties among agents with the same score, we randomly select T agents from among the tied nodes. We compute the betweenness centrality score for each agent following the algorithm proposed by (Brandes, 2001). Betweenness centrality measures the extent to which a node lies on the shortest paths between other nodes in the network (Latora et al., 2017). Nodes with high betweenness centrality are considered to be mediators (Zhang and Luo, 2017) and are often located on important bridges in the network, making them key players for information

flow and communication. Targeting these nodes can potentially have a greater impact on the overall network dynamics, creating barriers for anti-vaccine communities and preventing them from merging together. We explored additional centrality measures but found very limited differences (see Fig. A.1 in Appendix A for a comparison). For this reason, in the remainder of this chapter we will only consider betweenness centrality.

3.3.2.2 Dynamic Campaigns

Dynamic campaign strategies rely on the local information about individuals' vaccine opinions. Targeting based on neighborhood information has previously been utilized in other fields, including cooperation in multi-agent systems Lynch et al. (2018) and minimizing negative diffusion (e.g., Shi et al. (2019)). The following strategies outline the criteria for selecting the target set based on this information.

1. **Dynamic random strategy (DynRandT)**: at each time interval t_r, the target set is replaced with a new target set by selecting from the remaining population of agents with neutral opinions at random. We consider this campaigning strategy as a reference to evaluate the effectiveness of other dynamic selection criteria.
2. **Dynamic local information based strategies**: in this approach, our objective is to focus on agents with neutral opinions who are susceptible to negative influence from their social connections, i.e., agents with neutral opinion who have at least one anti-vaccine neighbour. This is a neighborhood-based scheme with two primary considerations: first, placing seeds to effectively inhibit the growth of the negative cluster, which requires them to have a certain number of negative neighbors; second, situating positive seeds in a way that maximizes the potential for positive clusters to grow (and eventually block negative clusters), which means targeting neutral agents with the greatest number of neutral neighbors. We expect a trade-off here. If agents have too many negative neighbors, they may become overwhelmed quickly, and thus positive influence might be wasted. Conversely, targeting agents with too many neutral neighbors might place them too far from negative clusters, thus becoming inefficient at blocking negative clusters from growing. We formalise this trade-off as follows:
   (a) **Local information based (DynAntiT)**: at each time interval t_r, the target set is updated and replaced with a new target set by selecting at random from the remaining population of agents with neutral opinions who have at least one anti-vaccine neighbor.
   (b) **Advanced Locl-Info with single-objective (DynLocT):** Here, we aim to target neutral agents who are in neighbourhoods that meet a trade-off

between blocking negative influence and allowing positive influence to spread. The first is related to the number of adjacent negative agents and the second relates to the number of adjacent neutral agents. Let ζ be the target number of anti-vaccine neighbours; here, we seek to target neutral agents who have a number of anti-vaccine neighbors as close as possible to ζ. In more detail, we do this by scoring agents according to the difference in their number of anti-vaccine neighbours from ζ as follows:

$$g_i(\zeta) = |n_i^- - \zeta|, \tag{3.2}$$

where n_i^- denotes the actual count of anti-vaccine neighbours for an agent i. Then, we select the T agents with the lowest score (and selecting at random in case of ties). If the number of agents selected is less than T, the selection process continues by choosing from the remaining population of agents with neutral opinions at random until T agents have been selected. We will vary ζ to identify the heuristic that best suppresses disease outbreaks.

(c) **Advanced local-info with multi-objective (DynAdvLocT)**: This heuristic builds on the previous heuristic, but we now also include the potential for positive information to spread by including the number of neutral neighbours in the scoring process. Again, let ζ be the target number of anti-vaccine neighbours and Z be the target number of neutral neighbours of an agent. Next, presuming that agent i has n_i^- anti-vaccine neighbours and n_i^0 neutral neighbours, we calculate a score according to:

$$g_i(\zeta, Z) = |n_i^- - \zeta| + |n_i^0 - Z|, \tag{3.3}$$

and select the T agents with the lowest scores. In the case of ties among agents with the same score, we randomly select T agents from among the tied nodes. If the number of agents selected is less than T, the selection process continues by choosing from the remaining population of agents with neutral opinions at random until T agents have been selected. Here, each term contributes independently to an agent's score and the additive structure allows for trade-offs between the two components, enabling control over whether the strategy prioritizes blocking negative diffusion or maximizing positive diffusion. In particular, by choosing a higher target number of neutral neighbors Z than the target number of anti-vaccine neighbors ζ, the campaign prioritizes individuals with more neutral neighbors and vice versa. As we will see below, the additive structure provides good performance. Other formulations may also be possible and could be explored in future work. In the following, we explore the dependence of the effectiveness of the heuristic on both target numbers of negative neighbours ζ and neutral neighbours Z.

3.4 Results

In this section, we present the obtained epidemic size for the proposed positive campaigns. During the first stage, i.e., opinion exchanges, we consider two cases as stopping criteria: one where opinions spread until all agents have adopted an opinion, referred to as $\tau = \infty$, which represents the long-run scenario and enables the evaluation of the long-term behavior. The other scenario where opinions spread over a certain period of time τ, represents the short-run scenario. In the short-run case, we compared the obtained epidemic size with that from a scenario where only anti-vaccine opinion diffusion is considered, previously investigated in Campbell and Salathé (2013). While our experimental setting differs from that of Campbell and Salathé (2013), we have applied their work within our setting for comparison purposes. We conducted extensive experiments to investigate the factors that determine the efficacy of each campaign in promoting vaccination.

Unless otherwise stated, the results show the epidemic size as a function of the social contagion rate parameter ω, where $\omega^- = \omega^+$, to compare the varying strengths of social influence on vaccine decision-making and their impact on the emergence of anti-vaccine communities and consequently disease spread. The social rate is a crucial component because it controls the growth of anti- and pro-vaccine communities. A low social rate implies that individuals are barely influenced by their social contacts, resulting in a low growth of anti-vaccine communities. In contrast, a high social rate indicates that individuals are highly influenced by their social contacts, leading to a large growth of anti-vaccine communities. To simplify the analysis, we have assumed equal social rates for pro- and anti-vaccine diffusion and evaluated the impact of different rates of positive general exposure μ^+ on the epidemic size, while keeping the negative general exposure μ^- fixed. Fixing μ^- allows for a clearer interpretation of the effects of the campaign budget relative to the strength of the negative influence, as the primary consideration is the relationship between μ^- and the counter-influence rate μ^+, specifically whether they exert influence at similar rates or if one is dominant over the other. The shaded area and error bars represents the 95% confidence intervals. Due to the small error values and the large y-axis scale, the shaded confidence interval area might not be clearly visible in some figures.

In all experiments, following a configuration similar to that used in Campbell and Salathé (2013), we consider a small-world network with size $N = 5000$, rewiring probability $p = 0.01$, average degree $\langle k \rangle = 10$, and the social rate range $\omega^- = [10^{-4}, 10^{-2}]$. This network size is adequate to capture the general trends of the system dynamics, namely information diffusion and disease transmission. Similarly, an average degree of 10 is consistent with ranges observed in empirical studies of regular social contacts (Read et al., 2008) and provides a balance between sparse and dense connectivity. Additionally, we use an opinion formation threshold of $\theta = 2$ as

we consider complex contagion in which agents require more than one exposure to adopt an opinion. For the SIR parameters, also following Campbell and Salathé (2013) as a baseline, we set the infection rate $\beta = 0.1$, recovery rate $\gamma = 0.1$, and seed set $I_0 = 1$. For the general negative exposure rate μ^-, we consider a value of 0.001, chosen to be close to the positive influence μ^+. For the additional parameters introduced in this study, specifically (ω^+, T, t_r, ζ, Z), we performed parameter sweeps for each. Unless otherwise stated, for each scenario we generate 500 different networks and for each network we run 500 SIR infection simulations.

3.4.1 Results of the Random Campaign

This section presents the epidemic size obtained by applying the random campaign StatRandAll to disseminate positive vaccine-related information, as illustrated in Fig.3.3.

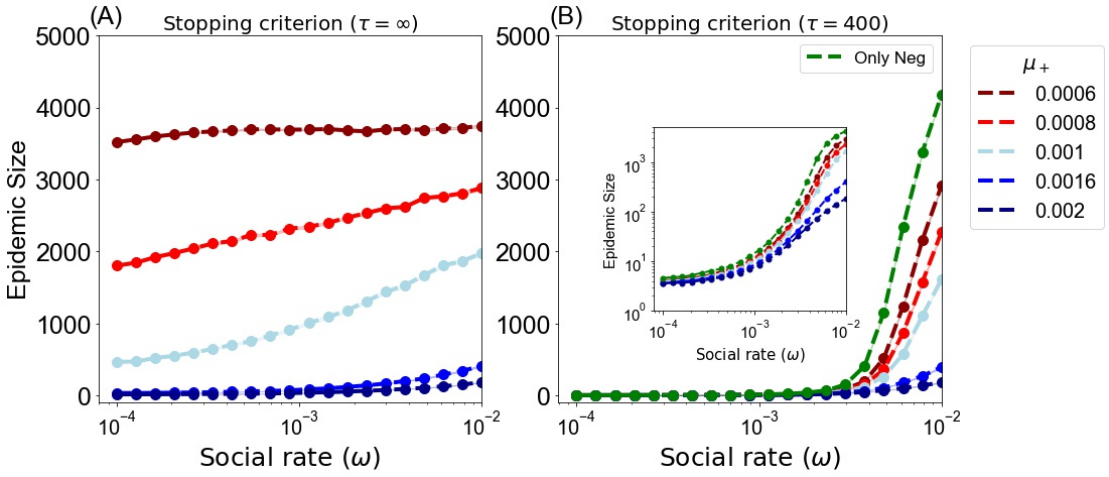

FIGURE 3.3: Average epidemic size for the random campaign (StatRandAll) as a function of the social rate $\omega = \omega^+ = \omega^-$. (A) $\tau = \infty$ (B) $\tau = 400$. The figure shows results for different positive exposure rates μ^+ with fixed negative exposure rate μ^- = 0.001.

Fig.3.3 (A) illustrates the long-run scenario with $\tau = \infty$ for varying social rates. This campaign seeds positive influence widely by targeting the whole population at random, resulting in a large number of seeds being generated, with the same occurring for negative seeds. When the social rate ω is low (e.g., $\omega < 10^{-3}$ in Fig.3.3(A)), the social influence of these seeds is slow, resulting in unconnected and smaller homogeneous communities, namely anti-vaccine and pro-vaccine communities. In addition, the existence of pro-vaccine communities prevents the merging of anti-vaccine communities. This ultimately leads to a smaller epidemic size compared to a higher social rate.

In contrast, when the social rate is high (e.g., $\omega = 10^{-2}$ in Fig.3.3 (A)), the social influence of these seeds spreads rapidly, leading to fast expansion of the communities

and allowing merging of communities. In this scenario, due to the rapid diffusion of opinions, the positive campaign fails to exert more influence over time, as the majority of individuals have already adopted an opinion. This ultimately results in large-sized communities and subsequently a higher epidemic size within anti-vaccine communities. This pattern is observable when the negative and positive general exposure rates diffuse at nearly equal rates, represented by red and light-blue lines in the figure. Nevertheless, with a much lower positive general exposure rate $\mu^+ \ll \mu^-$ (dark-red line), this pattern is almost nonexistent, since the positive seeding rate is low, rendering the negative influence dominant regardless of the social rate. In contrast, when the positive rate is much greater than the negative rate $\mu^+ \gg \mu^-$, represented by blue and dark-blue lines, the epidemic size dramatically decreases to less than 50 at the lowest social rate, i.e., $\omega = 10^{-4}$. Although the social rate does not play a significant role in such scenarios, higher social rates, namely $\omega = 10^{-2}$, lead to an increase in the epidemic size to less than 500.

Moreover, an increase in the positive external rate μ^+ results in a decrease in the epidemic size. The disparity in budget allocations for negative and positive external rates plays a pivotal role in shaping the prevalence of anti-vaccine opinions. A higher positive external rate than negative external rate leads to a dominance of positive influence, resulting in a smaller number of anti-vaccine opinions and consequently a smaller epidemic size, and vice versa.

Fig.3.3 (B) gives results for $\tau = 400$, and allows for a comparison between scenarios with and without positive campaigns. The figure distinctly illustrates that the propagation of positive vaccine-related information always yields a positive effect, even with a small positive budget (e.g., $\mu^+ = 0.0006$), resulting in a reduction in epidemic size compared to the scenario in which only anti-vaccine opinions are being spread (green line). Furthermore, the suppression of the epidemic increases as the positive general exposure rate μ^+ increases. However, as the social rate increases, the growth of communities also increases, leading to a higher epidemic size.

3.4.2 Results of the Static Campaigns

This section presents the epidemic size obtained by using targeted static campaigns to disseminate positive vaccine-related information. This includes the targeted random campaign StatRandT and the targeted centrality-based campaign StatCentT.

Fig.3.4 (A) and Fig.3.4 (C) illustrate the long-run scenario with $\tau = \infty$ for the targeted random and centrality-based campaigns, respectively. In general, the centrality-based campaign performs better in reducing the epidemic size than the targeted random campaign. Additionally, as the positive rate μ^+ increases, the performance of the campaigns improves, leading to higher epidemic suppression. However, for the

centrality-based campaign, this improvement is relatively small due to the fact that this campaign is efficient even when positive exposure rates are low.

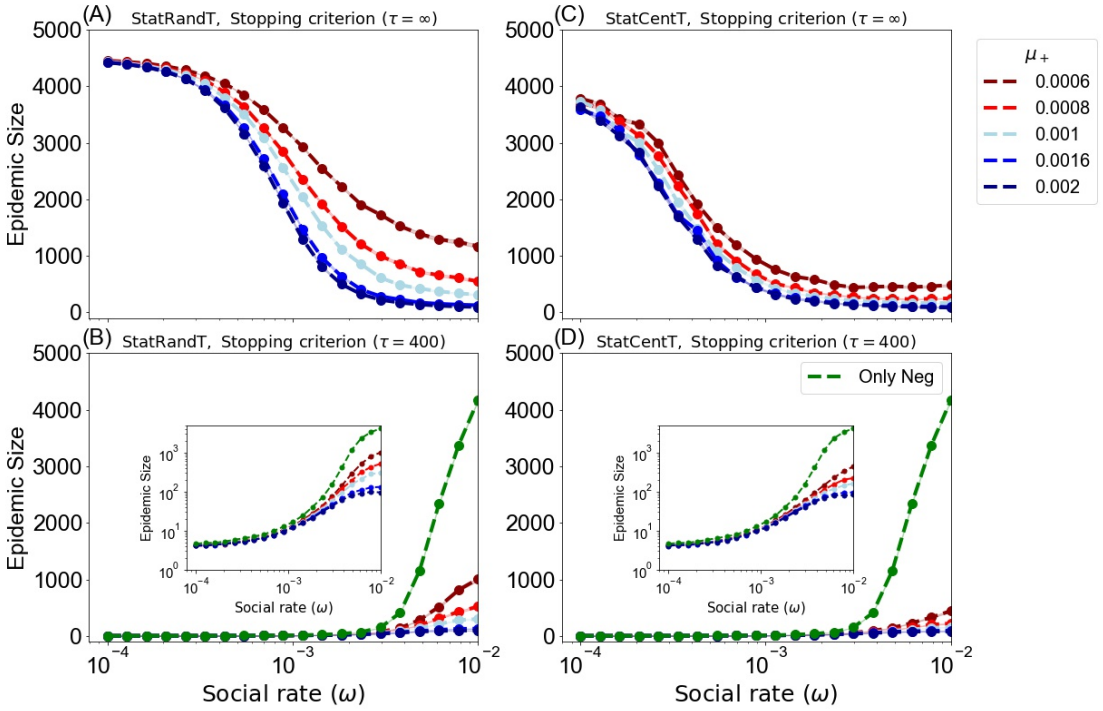

FIGURE 3.4: Average epidemic size for static campaigns as a function of the social rate $\omega = \omega^- = \omega^+$. (A) and (B) for the targeted random campaign (StatRandT). (C) and (D) for the centrality-based campaign (StatCentT). (A) and (C) $\tau = \infty$, (B) and (D) $\tau = 400$. The figures show different positive exposure rates μ^+ with fixed negative exposure rate $\mu^- = 0.001$. Target set size is $T = 500$. For each scenario we generate 300 different networks, and perform 300 SIR model runs for each network.

More importantly, there is a notable difference between the random (StatRandAll) and targeted campaigns in regard to the social influence rate ω. When the social rate is low, individuals tend to exchange opinions less frequently, and their vaccination behavior is largely influenced by the external campaigns. This is shown clearly in the long-run setting, where all individuals adopt a vaccine opinion, as seen in Figs 3.3 (A), 3.4 (A), and 3.4(C). Under these circumstances, the random campaign, see Fig.3.3 (A), tends to yield a smaller epidemic size compared to static campaigns, see Figs 3.4A, 3.4 (C). This is due to the fact that the random campaign generates a larger number of positive seeds over time than the targeted campaigns, which are restricted to a fixed set of agents. With lower levels of social interaction, the growth of homogeneous communities is slower, resulting in the formation of a large number of small, unconnected communities. More importantly, the scattered spread of positive seeds prevents mergers between the anti-vaccine communities. As a consequence, these smaller communities yield a smaller epidemic size.

In contrast, targeted campaigns generate limited seeds as they work with a specific and static target set and not the entire population, so their impact is restricted to the

positions of these seeds. In the centrality-based campaign, these positions are the most central nodes in the network, and as a result, they behave better in reducing the epidemic as they efficiently mitigate the connectivity between anti-vaccine communities compared to the targeted random campaign, where the target set is chosen at random.

On the other hand, when the social rate is high, targeted campaigns work better in containing the disease dynamics, resulting in smaller epidemic size as the strategically positioned targets reduce the connectivity between the anti-vaccine communities.

To further explain, consider the scenario where $\mu^- = \mu^+$, represented by light-blue curves. In the random campaign shown in Fig.3.3 (A), the epidemic size increases as the social rate ω increases. However, the targeted random (StatRandT) and targeted central campaigns (StatCentT) shown in Fig.3.4 (A) and Fig.3.4 (C), respectively, exhibit an inverse behavior, with the epidemic size decreasing as the social rate increases. The reason for this is that, at low social rates, broad seeding using the StatRandAll campaign distributes pro-vaccine seeds across the entire network, creating obstacles that efficiently mitigating the merging of emerging anti-vaccine communities. However, with a limited number of targets in the targeted campaign and a low rate of social interaction, these campaigns fail to mitigate the propagation of anti-vaccine influence, which is reinforced by broad and continuous exposures, as their effect only associates with their positions. On the other hand, at higher social rates, opinions are diffused faster, making it challenging to exert sustained influence over time. Thus, the StatRandAll campaign becomes less efficient compared to the targeted campaign, where the fixed positions of the targets successfully impede the connectivity of the anti-vaccine communities. The same behavior is observed regardless of the size of the positive campaign budget μ^+.

Furthermore, Fig.3.4 (B) and Fig.3.4 (D) display the system at time $\tau = 400$, allowing for a comparison between the campaign and the anti-vaccine opinion only scenario for the targeted random and centrality-based campaigns, respectively. The figure distinctly illustrates that the propagation of positive vaccine-related information results in a positive effect. As demonstrated in the figures, this approach reduces the epidemic size when compared to the scenario where only anti-vaccine opinions exist. Furthermore, the suppression of the epidemic increases as the positive general exposure rate increases. However, as the social rate increases, the growth of communities also increases, leading to a higher epidemic size.

A noteworthy observation is that the centrality-based campaign (StatCentT) is more effective in reducing the epidemic size than random campaigns compared to the scenario where only anti-vaccine opinions are spread. Additionally, the targeted random campaign (StatRandT) is more effective than the random campaign (StatRandAll). For instance, consider the scenario where $\mu^+ = \mu^-$, as shown by the

light-blue lines in Fig.3.3 (B), Fig.3.4 (B), and Fig.3.4 (D). At the highest social rate, i.e., $\omega = 10^{-2}$, the epidemic size for the random (StatRandAll), targeted random (StatRandT), and centrality-based campaigns (StatCentT) is 1615 ± 27, 302 ± 13, and 161 ± 5, respectively, compared to 4174 ± 15 for the scenario where only anti-vaccine opinions are being spread.

3.4.3 Results for the Dynamic Campaigns

This section presents the epidemic size obtained by applying targeted dynamic campaigns to disseminate positive vaccine-related information. This includes the dynamic random (DynRandT) and dynamic local information-based (DynAntiT) campaigns.

Fig.3.5 displays the results obtained by these campaigns. We evaluate the campaigns for different update times and results are presented as a function of the target set update time interval t_r. Here if the target set is changed very often, since multiple exposures are required for adopting opinions, the likelihood of each individual agent being influenced is very low. Correspondingly, fairly small amounts of influence are spread over a large set of agents that are targeted at different times. In contrast, when leaving the target set unchanged for longer, agents in the target set can accumulate multiple exposures which might lead to opinion adoption. However, this also implies that influence is not spread very widely and occasionally agents who already hold an opinion might be targeted. From these considerations it becomes clear that there must be an optimal switching time which maximizes the effect of the positive campaign.

Fig.3.5 (A) and Fig.3.5 (C) illustrate the long-run scenario with $\tau = \infty$ for varying update time intervals for the dynamic random DynRandT and DynAntiT campaigns, respectively. The results demonstrate that an optimal update time interval exists at around $t_r = 20$. This is particularly obvious when the general positive exposure rate μ^+ is lower than the negative one $\mu^+ \ll \mu^-$. This optimal time results from the trade-off explained above. Therefore, with a lower general positive exposure rate, the campaign performs better with relatively slow updates.

On the other hand, as the general positive exposure rate increases, indicating a stronger positive influence, the pronounced effect—i.e., an optimal time at $t_r = 20$—diminishes. This effect becomes almost negligible when $\mu^+ > \mu^-$ in the DynRandT campaign, and when $\mu^+ \geq \mu^-$ in the DynAntiT campaign. In such cases, the fastest update strategy $t_r = 1$ becomes an effective option. This phenomenon occurs because a higher general positive exposure rate increases the probability of positive influence, and when combined with quick updates, it allows us to attain widespread coverage by targeting susceptible agents before they are negatively influenced. Moreover, after a time interval of $t_r = 10$, the epidemic size continues to

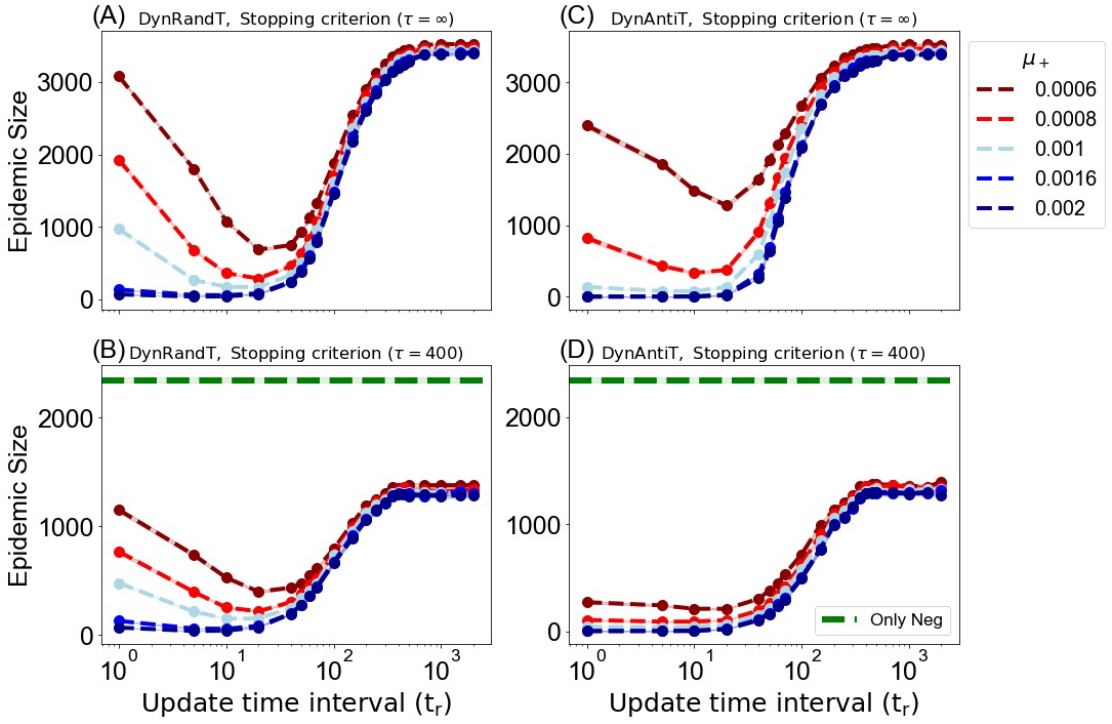

FIGURE 3.5: Dependence of the average epidemic size on the campaign updating interval using dynamic campaigns.(A) and (B) represent the dynamic random campaign DynRandT, and (C) and (D) represent the DynAntiT campaign. (A) and (C) $\tau = \infty$, (B) and (D) $\tau = 400$. The targets set $T = 50$, social rate is $\omega^+ = \omega^- = 0.006$. The figures show different positive exposure rates μ^+ with fixed negative exposure rate $\mu^- = 0.001$.

increase as the interval t_r increases until reaching a stationary state, where there are no further changes in the epidemic size. As we increase the update intervals, we reduce the scope of our targeting coverage, resulting in a corresponding decrease of the positive effects we initially achieved in mitigating the negative influence.

Furthermore, comparing the DynRandT and DynAntiT campaigns, Fig.3.5 (A) and Fig.3.5 (C) demonstrate varying performance in reducing the spread of an epidemic. For the smallest positive general exposure rates displayed in the figure, e.g., dark-red lines and $\mu^+ = 0.0006$, which is much smaller than the negative general exposure rate $\mu^- = 0.001$, the dynamic random strategy outperforms the DynAntiT strategy at the optimal time $t_r = 20$, resulting in a smaller epidemic size. At time $t_r = 20$, the epidemic size is 1276 ± 29 and 688 ± 19 with $\mu^+ = 0.0006$, and 375 ± 13 and 288 ± 8 with $\mu^+ = 0.0008$ for DynAntiT and dynamic random DynRandT, respectively. However, as the general positive exposure rate increases, the DynAntiT strategy is more effective at mitigating the negative influence. This is also noticeable when both positive and negative rates spread at the same rate, i.e., $\mu^+ = 0.001$, where epidemic size at time interval $t_r = 1$, is 139 ± 8 and 972 ± 23 for DynAntiT and DynRandT, respectively. Furthermore, the epidemic size reduced even further when the positive general exposure rate was much greater than the negative rate with $\mu^+ = 0.002$ and

$\mu^- = 0.001$ and reached 3 ± 0.08 and 70 ± 2 for the DynAntiT and DynRandT campaigns respectively, at $t_r = 1$.

Fig.3.5 (B) and Fig.3.5 (D) display the system at time $\tau = 400$ for the dynamic random and DynAntiT campaigns, respectively, allowing for a comparison between the campaigns and the only anti-vaccine opinion scenario. Once again, the dissemination of positive vaccine-related information yields a positive impact in mitigating the diffusion of anti-vaccine sentiments compared to the scenario where only anti-vaccine opinions are disseminated. Furthermore, the suppression of the epidemic increases as the positive general exposure rate increases.

3.4.4 Results for the Advanced Local-Info Campaigns

This section presents the epidemic size obtained when applying the advanced targeted dynamic campaigns to disseminate positive vaccine-related information. This includes the targeted advanced local-info (DynLocT) and multi-objective advanced Locl-Info (DynAdvLocT) campaigns. For the DynLocT campaign, the results are presented as a function of the target number of anti-vaccine neighbors ζ to assess the impact of this parameter on mitigating the growth of anti-vaccine communities and, consequently, reducing the size of the epidemic. For the DynAdvLocT campaign, the results are presented as a function of both ζ and Z, with different values of μ^+ to explore the trade-off between maximizing the growth of pro-vaccine communities and minimizing the growth of anti-vaccine communities. The results are shown in Fig.3.6 and Fig.3.7 for DynLocT and DynAdvLocT campaigns, respectively.

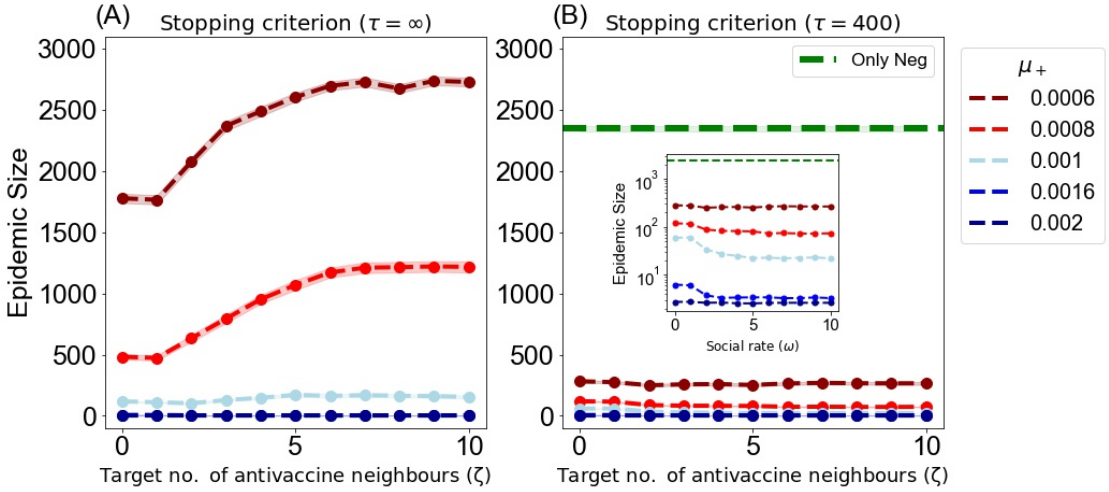

FIGURE 3.6: Dependence of the average epidemic size on the target number of anti-vaccine neighbors ζ using the DynLocT campaign. The updating time is $t_r = 1$, the social rate is $\omega = 0.006$ for both negative and positive $\omega^+ = \omega^-$, and the size of the target set is $T = 50$. The figures show different positive exposure rates μ^+ with fixed negative exposure rate $\mu^+ = 0.001$.

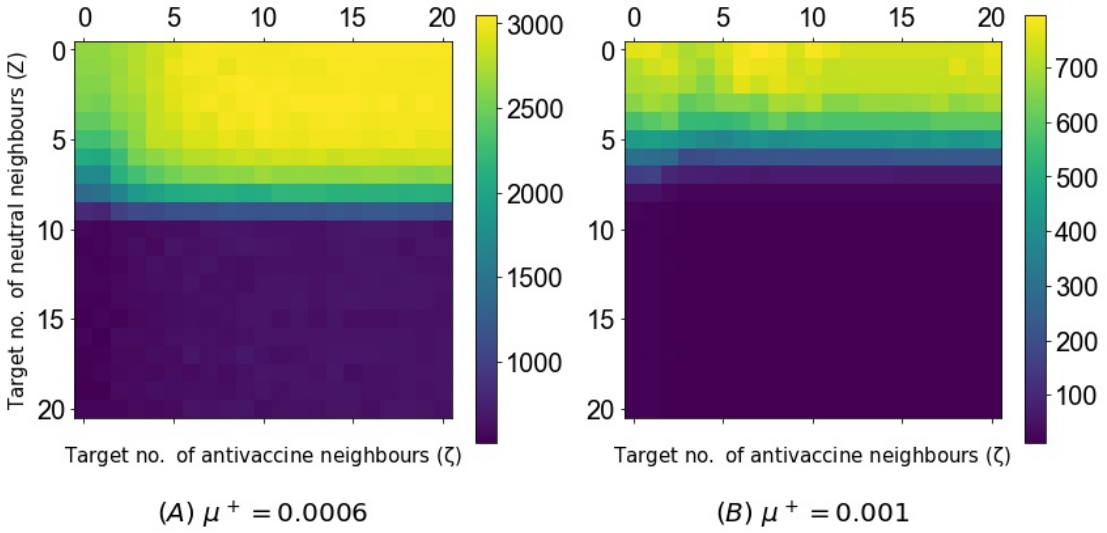

FIGURE 3.7: Average epidemic size using the DynAdvLocT dynamic campaign in the long-run setting of $\tau = \infty$. The figure shows the performance of the dynamic campaigns with $T = 50$ targets. The epidemic size is shown as a function of the target number of anti-vaccine neighbors ζ and the target number of neutral neighbors Z a neutral has at time t. The updating time is t_r =1, and the social rate is $\omega = 0.006$ for both negative and positive $\omega^+ = \omega^-$. The general exposure influence rate for negative is $\mu^- = 0.001$ and the positive rate is shown in the figures captions.

For DynLocT campaign, considering the long-run behavior with $\tau = \infty$, see Fig.3.6 (A), we found that optimal performance is obtained for $\zeta \leq 1$, which is particularly obvious when the positive rates μ^+ are significantly lower than the negative rates μ^-, as indicated by the dark-red line in the graph. We also observed the same behavior found in the dynamic campaigns, where fast updates with weak influence rates are not effective in convincing the target set before they are updated since agents require multiple exposures to adopt a particular opinion. The challenge becomes greater when selecting neutral agents who are surrounded by many anti-vaccine adopters, as they are more likely to be negatively influenced by their social contacts. However, this behavior diminishes as the positive influence rate increases and leads to a significant reduction in the epidemic size, particularly when $\mu^+ \gg \mu^-$, regardless of the ζ value. For example, the epidemic size remains the same at $\zeta = 1$ and $\zeta = 8$, taking a value of 3 ± 0.1, see Fig.3.6A at μ^+ = 0.002.

In the case of the DynAdvLocT campaign, similar to DynLocT, we have observed that an optimal reduction in the epidemic size is obtained for $\zeta \leq 1$, particularly when the positive rates μ^+ are considerably lower than the negative rates μ^-, as depicted in Fig.3.7 (A). Moreover, this reduction increases as Z increases, indicating a focus on targeting individuals with a larger number of neutral neighbors. However, there is a remarkable observation that when $Z \geq 10$, the epidemic size is significantly diminished irrespective of the ζ value, which is equivalent to the average degree of the network $\langle k \rangle = 10$. Furthermore, when $Z > 10$, the epidemic size remains relatively constant.

In practice, this corresponds to prioritizing individuals with few negative neighbors and a large number of undecided contacts who have not yet formed a vaccine stance and are more likely to be positively influenced by the targeted agent. This is achieved by choosing larger values of Z and smaller values of ζ, which leads to targeting individuals with a higher number of neutral neighbors. Larger values of ζ indicate prioritizing individuals surrounded by many negative neighbors. Targeting such individuals therefore is less effective as the negative influence is stronger than the campaign's impact, and the dissemination of positive influence becomes more challenging in communities already saturated with negative influence. This consequently leads to poorer control of negative propagation and larger epidemic sizes.

A similar behavior is observed in the case of $\mu^+ = \mu^-$, as shown in Fig.3.7 (B), where a more significant reduction in the epidemic size is achieved with fewer target number of neutral neighbours Z compared to the previous scenario. The figure indicates that we can achieve a greater reduction when $Z > 7$. As an example, this campaign efficiently reduced the epidemic size to 12 ± 0.5 when $\zeta = Z = 10$. Furthermore, beyond this point, this reduction remains relatively constant, and increasing the Z value does not yield any additional effect. In this scenario, the number of anti-vaccine neighbors, does not have a major effect.

To elucidate the targeting strategy of this campaign, we illustrate in Fig.3.8 the neighborhood structure for the targeted agents of this campaign in conjunction with the evolution of opinion diffusion for both anti- and pro-vaccine adopters. In this campaign, by choosing a higher target number of neutral neighbors Z than the target number of anti-vaccine neighbors ζ, we prioritize the number of neutral neighbors over the number of anti-vaccine neighbors and vice versa. In the former scenario, this strategic shift directs positive allocation to agents more likely to propagate positive influence to a greater extent, as they have a larger number of neutral neighbors than anti-vaccine neighbors. This potentially will maximize the growth of pro-vaccine communities while restricting the growth of the anti-vaccine communities by being in proximity to them. In Fig.3.8 (A), an illustrative example of such a scenario is presented, with $Z = 8$ and $\zeta = 1$. As observed in the figure, this scenario focuses on targeting agents with the highest number of neutral neighbors and the smallest number of anti-vaccine neighbors. The corresponding evolution of opinions, depicted in Fig.3.8 (D) and (E), demonstrates that this targeting scheme maximizes the number of pro-vaccine adopters while simultaneously minimizing the number of anti-vaccine adopters, as indicated by the green lines. Consequently, it leads to a greater reduction in the epidemic size, as depicted in Fig.3.8 (F) with the green bar.

Conversely, when the target number Z is smaller than ζ, we prioritize agents with a higher number of anti-vaccine neighbors than neutral neighbors. In this instance, we shift the targeting to focus on agents with a higher number of anti-vaccine neighbors

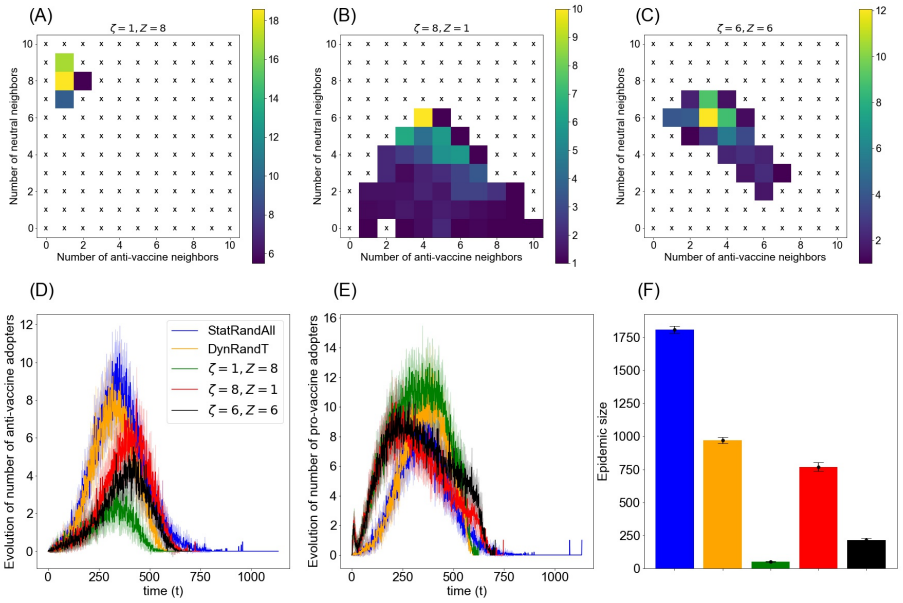

FIGURE 3.8: Targeting scheme for DynAdvLocT campaign. The top panels illustrate the neighborhood structure of the target set at a single time step t=350 during the opinion diffusion stage, specifically showing the number of anti-vaccine neighbors and neutral neighbors. These panels represent the average number of agents with x anti-vaccine neighbors and y neutral neighbors for various settings: (A) $\zeta = 1, Z = 8$, (B) $\zeta = 8, Z = 1$, (C) $\zeta = 6, Z = 6$. (D) and (E) illustrate the evolution of anti-vaccine opinion adopters and pro-vaccine opinion adopters, respectively, while panel (F) represents the corresponding epidemic size. 'x' indicates that no agent exists for that neighborhood pattern. The targeting analysis is an average of 15 simulations.

and a relatively smaller number of neutral neighbors, see Fig. 3.8(B) for an example with $Z = 1$ and $\zeta = 8$. Although emphasizing agents with more anti-vaccine neighbors can protect them from the anti-vaccine influence, it is less efficient in mitigating the overall propagation of this influence in the network. As depicted in Fig. 3.8(D) and (E), this targeting scheme results in a higher number of anti-vaccine adopters and a lower number of pro-vaccine adopters compared to the first scenario (i.e., $Z = 1$ and $\zeta = 8$), as seen in the green and red lines. Correspondingly, it leads to a higher epidemic size, even though we are protecting the most vulnerable agents.

Furthermore, in Fig.3.8(C), we depict a scenario that assigns high priority to both Z and ζ, where $Z = \zeta = 6$. In this scenario, as observed, we target agents with a relatively high number of neutral neighbors and a high number of anti-vaccine neighbors, avoiding the smaller numbers of neutral neighbors targeted by the scenario $Z = 1, \zeta = 8$. Consequently, this approach proves more efficient than the scenario depicted in Fig.3.8 (B), resulting in lower numbers of anti-vaccine adopters, as depicted in the black line in Fig.3.8 (D), and accordingly, a smaller epidemic size in Fig.3.8 (F) with the black bar. Despite this improvement over the previous scenario,

the first scenario prioritizing neutral neighbors (Fig.3.8 (A)) yields the best mitigation of negative influence and reduction in epidemic size. We further compare all three scenarios to random and dynamic random campaigns, depicted by the blue and orange colors, respectively, in opinion evolution and the corresponding epidemic size. As demonstrated in the figure, the DynAdvLocT scenario outperforms the benchmark cases in all scenarios.

3.4.5 The Impact of the Target Set Size

The size of the target set is a crucial factor in determining the effectiveness of the targeted campaigns. Figure 3.9 demonstrates the correlation between the size of the target set and the extent of the epidemic in targeted campaigns. In the case of static campaigns, see Fig.3.9 (A) and Fig.3.9 (B), a larger target set results in better mitigation of anti-vaccine diffusion than a small target set. This is due to the positive allocation sticking around only these agents, and a relatively larger set ensuring a fair coverage in the network. However, we observed that a very large target set might act as noise and impede the focus of the campaign. For example, in the centrality-based campaign, after a certain point, i.e., T=500 which represents 10% of the population, the curve starts to increase again due to the inclusion of numerous low centrality agents in the target set, which hampers the centrality effects and introduces more randomness in the selection process.

In dynamic campaigns, on the other hand, the results have shown that smaller target group sizes, particularly in DynAntiT, DynLocT, and DynAdvLocT, yield better reduction compared to larger sizes, as depicted in Fig.3.9(D-F). The dynamic random campaign DynRandT demonstrates a slight increase as the target size increases. These findings suggest that it is more effective to direct resources to a smaller yet changeable target set. This observation can be attributed to the fact that a smaller size allows for more selective targeting of individuals who meet the campaign's criteria and are more likely to be influenced. On the contrary, a larger target set may include individuals who are less susceptible to the negative influence, introducing more randomness in the selection process and resulting in decreased effectiveness. Furthermore, the positive strength allocated to a campaign is typically distributed evenly across the target group. As a result, a change in the target size can lead to a change in the strength allocated per individual.

3.4.6 Cross-Campaign Comparison

In this section, we present a comparison of the proposed campaigns to evaluate their performance in reducing the epidemic size. Further performance evaluation of the campaigns' impact on disease dynamics during the course of the epidemic can be seen

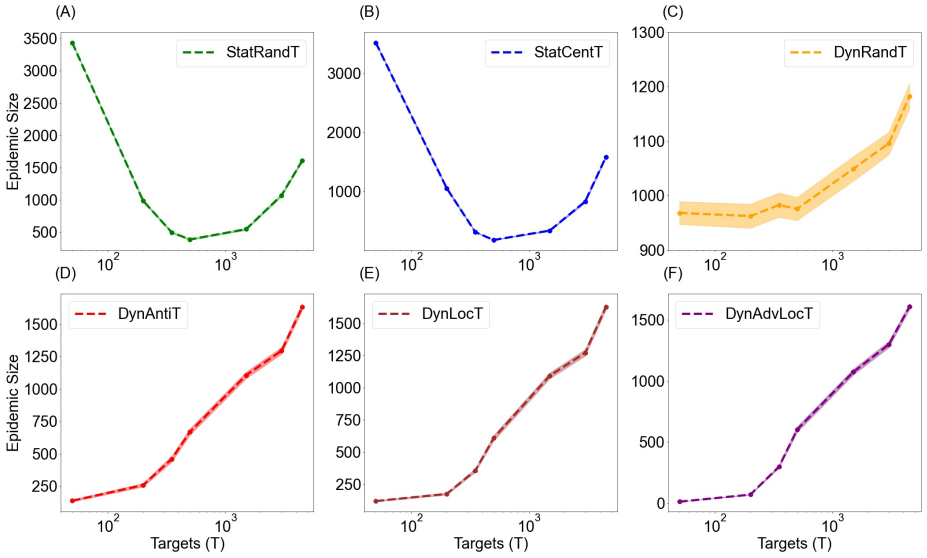

FIGURE 3.9: Epidemic size obtained with varying target sizes for targeted campaigns in the long-run setting where $\tau = \infty$. The figure illustrates the epidemic size obtained for all campaigns with general exposure rates $\mu^- = \mu_+ = 0.001$. The social rate is $\omega^+ = \omega^+ = 0.006$ for all scenarios. For dynamic campaigns, the updating time interval is t_r =1, for DynLocT $\zeta = 1$, and DynAdvLocT campaigns the target numbers of negative and neutral neighbours are $\zeta = 10, Z = 10$. For each scenario we generate 500 different networks, and perform 500 SIR model runs for each network.

in Fig. A.2 in Appendix A. As we see in Fig. A.2 in Appendix A, our campaign achieves a reduction in peak numbers of infections proportionate to the sizes of the outbreaks. In Fig. 3.10, we compare the best scenarios for each campaign, depicting three distinct states of strength allocation in the general exposure for negative and positive campaigns, considering a high social rate—a crucial scenario in which anti-vaccine communities can expand significantly. The aim is to investigate which campaign can mitigate this expansion most effectively. The figure also demonstrates the long-run behavior of the system. Across all scenarios, advanced local information campaigns, particularly DynAdvLocT , achieve the best performance in reducing the epidemic size.

In situations where the positive general exposure rate is much smaller than the negative general exposure rate (i.e., $\mu^+ \ll \mu^-$), the best practice is to use DynAdvLocT with slower updates (i.e., $t_r = 20$). The second-best option is the centrality-based campaign. The former results in an epidemic size of 345 ± 13, while the latter results in 493 ± 15.

When both negative and positive general exposures exert the same rate of influence (i.e., $\mu^+ = \mu^-$), DynAdvLocT with fast updates (i.e., $t_r = 1$) is the most effective, while centrality-based and other dynamic campaigns produce relatively similar performance. Finally, when the positive general exposure rate is higher than the

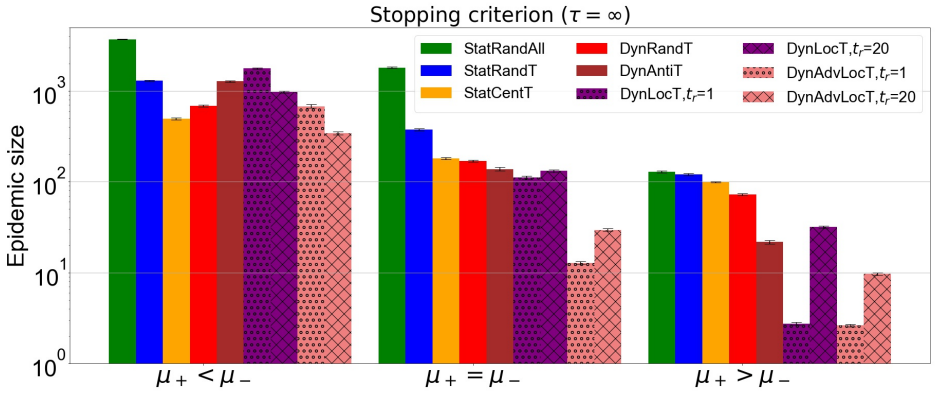

FIGURE 3.10: Cross-campaign epidemic size comparison with $\tau = \infty$. The figure illustrates the epidemic size obtained for all campaigns with different positive rates μ^+ compared to negative rate μ^-. For all campaigns, $\mu^- = 0.001$. First group: $\mu^+ = 0.0006$, second group: $\mu^+ = 0.001$, and third group: $\mu^+ = 0.002$. Social rate is $\omega = 0.006$ for all scenarios. Target set size $T = 500$ for static campaigns, i.e., StatRandT, and StatCentT, and $T = 50$ for the other dynamic campaigns. For dynamic campaigns, the updating time is $t_r = 20$, in addition we include $t_r = 1$ for DynLocT and DynAdvLocT campaigns. In addition, for DynLocT ζ =1, for DynAdvLocT $\zeta = Z$ =10.

negative one (i.e., $\mu^+ > \mu^-$), the best practice is to use either DynLocT or DynAdvLocT with fast updates (i.e., $t_r = 1$).

From Fig.3.10, it is observed that although the random strategy StatRandAll (green bars) reduces epidemic size, it is less effective than other campaigns, resulting in higher a epidemic size. The static random campaign StatRandT shown by blue bars follows a similar pattern but performs better than the StatRandAll campaign. This observation, as investigated earlier, is due to the fact that the random campaign StatRandAll fails to exert more influence over time, and opinions become socially-driven. In contrast, the targeted random campaign StatRandT targets fixed positions in the network, creating barriers that prevent the clustering of anti-vaccine communities, making it slightly more effective than the random campaign StatRandAll. This observation is consistent with the study conducted by Zhang et al. (2015), where the authors demonstrated that a random selection of the seed set resulted in poor mitigation of misinformation propagation.

In Figure 3.11, we consider a scenario where negative information spreads faster than positive information, i.e., $\omega^- > \omega^+$. This is motivated by existing literature, which highlights distinct diffusion dynamics between positive and negative information, with negative information spreading more rapidly and extensively than positive information (Tsugawa and Ohsaki, 2015). This raises critical questions: Can the diffusion of negative information still be effectively controlled in such a context? Moreover, how do the proposed campaigns perform under these conditions? As shown in Fig.3.11, although the performance is slightly worse compared to the scenario with faster positive spreading depicted in Fig.3.10, it is still possible to

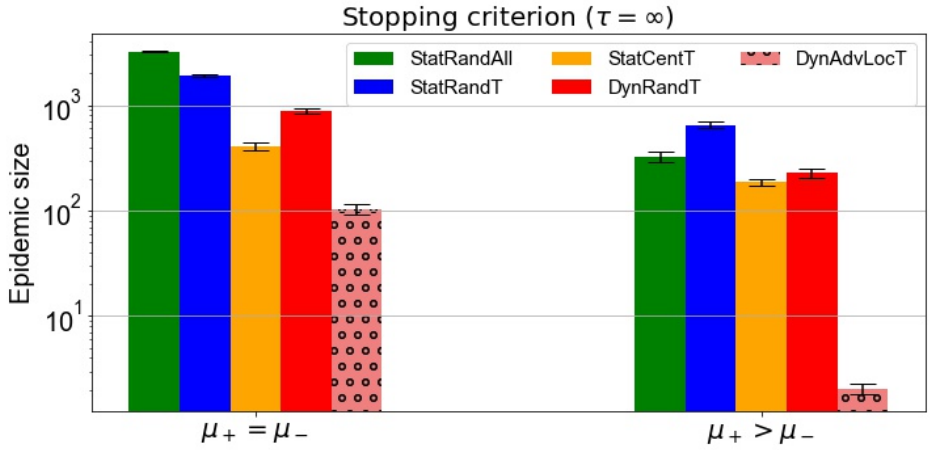

FIGURE 3.11: Cross-campaign epidemic size comparison with $\tau = \infty$. The figure illustrates the epidemic size obtained for the best campaigns with different positive rates μ^+ compared to negative rate μ^- and for lower positive social rate. For all campaigns, $\mu^- = 0.001$. First group: $\mu^+ = 0.001$ and second group: $\mu^+ = 0.002$. Social rate is $\omega^- = 0.006$, $\omega^+ = 0.003$ for all scenarios. Target set size $T = 500$ for static campaigns, i.e., StatRandT, and StatCentT, and $T = 50$ for the other dynamic campaigns. For dynamic campaigns, the updating time is $t_r = 20$ for DynRandT and $t_r = 1$ for DynAdvLocT. For DynAdvLocT $\zeta = Z$ =10. The results are the average of 50 simulations.

effectively contain the diffusion and maintain a low epidemic size, particularly with the use of the DynAdvLocT strategy. Furthermore, we also note that the performance improves significantly as the positive budget increases.

3.5 Summary

In this chapter, we proposed a model to explore effective methods for mitigating the spread of anti-vaccine opinions and reducing the size of epidemics by applying positive counter-campaigns that spread positive vaccine-related information. The model specifically captures scenarios relevant to childhood diseases, such as measles, where vaccination decisions occur during early childhood while outbreaks often occur later during school age. Although the real world is considerably more complex than our assumptions, the preliminary results presented here suggest that interventions aiming to mitigate anti-vaccine social contagion can help in controlling pathogen spread. We proposed efficient heuristics based on two main paradigms: social network structure and negative local information, in addition to two control schemes static and dynamic. Through extensive experiments, we systematically analyzed the performance of the proposed strategies, identifying their strengths and limitations. In particular, we examined how different campaigns affect the diffusion of anti-vaccine opinions, the connectivity among emerging anti-vaccine communities, and the resulting epidemic size.

We have demonstrated that the existence of positive influence propagation can impede the flow of anti-vaccine attitudes and mitigate the emergence of anti-vaccine communities, which helps suppress the spread of epidemics. However, this impact varies across different campaigning approaches. Specifically, we have demonstrated that targeted campaigns that select a subset of the population based on certain criteria have been found to be more effective in suppressing the epidemic compared to the random campaign. In addition, we find that the dynamic approach is more efficient than the static approach. Nonetheless, the effectiveness of these strategies varies based on several key factors: positive budget size, social influence strength, target set size, and the time horizons for dynamic campaigns.

Moreover, centrality-based and local information-based strategies have shown superior performance. The centrality-based campaign places positive seeds strategically on the most central bridges within the network, thereby preventing the merging of anti-vaccine communities more efficiently. However, this method requires a large number of targets to be effective, making it impractical in limited-resource settings. Local information-based strategies, on the other hand, prove to be the most effective, as they prevent the most vulnerable agents from being negatively influenced. By focusing on individuals who are more vulnerable to negative influences, this technique outperforms random selection and even outperforms the centralized approach. It also demonstrates efficacy even with a small target set. Quantifying the amount of negative influence on neutral social contacts yields further reductions in the size of the epidemic. Additionally, incorporating the potential for maximizing positive influence, by considering the number of neutrals for each candidate, leads to higher effectiveness, as seen in DynAdvLocT strategy. However, it is important to note that these local information-based campaigns assume complete knowledge of the vaccine-related attitudes within the population.

One crucial factor in the diffusion of anti-vaccine sentiments and the formation of anti-vaccine communities within a social network is the level of social influence between individuals; as the social contagion rate increases, the size of the anti-vaccine communities grows, consequently increasing the epidemic size. This observation is consistent with previous studies that investigated a similar problem and demonstrated the role of social interactions in promoting the growth of anti-vaccine communities, thus increasing the epidemic size (Campbell and Salathé, 2013; Dorso et al., 2017). Therefore, we found that targeted campaigns can effectively contain the spread of anti-vaccine diffusion in such a scenario by creating barriers that prevent the clustering of anti-vaccine communities, which helps reduce the final epidemic size. Additionally, the dynamic control approach, which involves continuously updating the target set, has been found to be more effective in suppressing the epidemic compared to a static control approach, as it provides continuous and iterative

exposure to positive messaging while keeping the campaign involved with the evolution of anti-vaccine attitudes.

In conclusion, this chapter investigated the impact of different strategic intervention campaigns aimed at combating the spread of anti-vaccine opinions on the resulting epidemic size. In the next chapter, we extend this investigation by leveraging optimization techniques to design new strategies aimed at minimizing the propagation of negative influence. We introduce a novel approach by formulating the problem as an optimization problem with two distinct objectives: (1) minimizing the spread of anti-vaccine influence, and (2) maximizing the spread of pro-vaccine influence. We address problem formulation, key challenges, solution optimality compared to the best-performing methods presented in this chapter, and the trade-offs between the two optimization goals.

Chapter 4

Optimization-based Campaigns for Vaccine Opinion Adoption

Parts of this chapter have been published in the proceedings of the International Conference on Complex Networks and Their Applications (Alahmadi et al., 2025a), and have been published as a journal paper in Applied Network Science (Alahmadi et al., 2025c).

4.1 Introduction

In this chapter, we extend our investigation into strategies for reducing epidemic outbreaks by influencing the social dynamics of vaccine-related opinions. Here, we introduce a novel approach by formulating the problem as an optimization problem that aims to minimize anti-vaccine opinions. To this end, we establish a mathematical framework and leverage optimization algorithms to design new heuristics to control opinion diffusion. By integrating optimization techniques with epidemiological models, we aim to provide a systematic approach for developing targeted and adaptable strategies that help policymakers implement efficient campaigns and improve public health outcomes.

Optimal control theory and optimization techniques have been widely employed in the existing literature on epidemic control (Kar and Jana, 2013; Nowzari et al., 2015; Mondal and Khajanchi, 2022; Doyle et al., 2024; Acemoglu et al., 2024; Sun et al., 2023; Federico et al., 2024). Specifically, researchers have investigated epidemic control through different control mechanisms, such as treatment, vaccination, and isolation (Kar and Jana, 2013; Mondal and Khajanchi, 2022; Sun et al., 2023; Federico et al., 2024). Other studies have explored epidemic control considering different perspectives, such as identifying cost-optimal strategies under budget constraints

(Nowzari et al., 2015), developing cost-effective lockdown strategies with vaccine uncertainty informed by policymakers' priorities (Doyle et al., 2024), and minimizing testing costs (Acemoglu et al., 2024). The majority of these studies address the problem by introducing control variables into compartmental epidemic models represented as systems of ordinary differential equations (ODEs). These include the classic SIR model (Acemoglu et al., 2024), the age-stratified SEIR model (Doyle et al., 2024), as well as variants of SIR such as SIRV (Kar and Jana, 2013), SAIQJR (Mondal and Khajanchi, 2022), and SIRS (Federico et al., 2024). Nevertheless, none of these studies investigate optimal epidemic control from the perspective of mitigating the spread of vaccine misinformation through network-targeted strategies and the optimal deployment of influence resources.

In addition to their application in controlling epidemics, optimization techniques have also been applied to manage the diffusion of information and opinions. For example, they have been used in the influence maximization problem (e.g., (Masucci and Silva, 2014; Lynn and Lee, 2016; Eshghi et al., 2018)), where the goal is to maximize the number of individuals adopting a specific opinion. Conversely, they have also been applied in misinformation mitigation, or the influence minimization problem, (e.g., (Budak et al., 2011; Zhang et al., 2015; Tong et al., 2018; Yang et al., 2019a)), where the primary objective is to minimize the spread of undesirable influence within a social network. Motivated by these approaches, we relate the epidemic control problem to the influence control problem (i.e., influence minimization and maximization) within our control framework, aiming to identify optimal control strategies for managing opinion diffusion and epidemic spread.

Establishing a direct analytical link between the final epidemic size and the influence exerted during the opinion dynamics phase is challenging. This difficulty arises from the complex and stochastic nature of individual interactions involved in both opinion diffusion and disease transmission, which makes it difficult to predict exactly how a specific influence strategy would translate into a specific reduction in epidemic size. However, for measles-like diseases where vaccination confers full immunity (CDC, 2025), disease transmission occurs primarily among unvaccinated individuals (i.e., anti-vaccine opinion adopters); consequently, the number of anti-vaccine agents is related to the epidemic size. In light of this correlation, we formulate the problem as a budget-constrained optimization problem aimed at controlling opinion diffusion dynamics, with the ultimate goal of reducing the final epidemic size. Specifically, we explore the impact of two optimization objectives on reducing the final epidemic size: (1) minimizing the number of anti-vaccine opinion adopters, (2) maximizing the number of pro-vaccine opinion adopters. While traditional influence minimization studies typically aim to reduce the number of individuals influenced by a specific opinion, and influence maximization focuses on increasing the number of adopters of a desired opinion, our objective goes beyond these goals. We are interested in the

distribution of anti-vaccine opinion adopters and aim to strategically disrupt their connectivity to prevent the formation of large, unprotected groups and ultimately reduce the final epidemic size. We explore the effect of these approaches in achieving this goal and, consequently, in reducing the final epidemic size.

Building on this foundation, we propose optimization intervention campaigns, designed to determine the optimal allocation of resources for targeting individuals. These campaigns are adaptive, continuously allocating resources in response to the evolving dynamics of anti-vaccine influence across the network. This approach is inspired by techniques from the influence maximization literature, particularly the continuous influence maximization method introduced by Romero Moreno et al. (2021). It contrasts with the common one-off targeting methods used in the field (Budak et al., 2011; Yang et al., 2019a; Fügenschuh and Fu, 2023) and the dynamic heuristics explored in the previous chapter. An overview of targeting methods for controlling information and opinion diffusion is provided in Section 2.5.1.

In this chapter, we make the following contributions to address research questions 1d and 2a detailed in Section 1.1. To begin with, we establish a mathematical framework that describes how opinions evolve in our model by deriving the equations of the opinion diffusion model described in Chapter 3. Then, we define two objective functions, formulate the optimization problems, and introduce the intervention campaigns. Next, we examine the performance of the campaign in controlling opinion diffusion and epidemic spread, and analyze the targeting strategy of these campaigns. Additionally, we evaluate the performance of the campaigns in a more realistic scenario where interventions are introduced at a later stage, after the spread of anti-vaccine opinions has already been established, to explore the impact of intervention timing on control effectiveness. This setting reflects real-world scenarios in which early intervention may not be feasible due to delayed awareness or slow policy response.

The chapter is structured as follows: Section 4.2 describes the opinion dynamics, formalizes the problems of minimizing the number of anti-vaccine agents and maximizing the number of pro-vaccine agents, and introduces the intervention campaigns. Section 4.3 presents the corresponding results: Section 4.3.1 shows the results of the minimization approach, while Section 4.3.2 presents the results of the maximization approach. Section 4.3.4 explores control performance in late intervention scenarios, and Section 4.3.5 examines the structure of anti-vaccine communities and its impact on the performance. Finally, Section 4.4 summarizes the main findings.

4.2 Model Description and Methods

This chapter builds upon the framework established in Chapter 3, maintaining the same diffusion models for opinions and epidemics as described in Section 3.2. In short, we consider a population of N agents interconnected via a small-world network, using the Watts-Strogatz small-world network model (Watts and Strogatz, 1998). Initially, we assume that all agents hold a neutral opinion about vaccines and can adopt either an anti-vaccine opinion, pro-vaccine opinion, or remain neutral. Motivated by the approach of Campbell and Salathé (2013), our model then consists of two diffusion stages: vaccine opinion diffusion followed by disease spread. Inspired by their work, we developed a model for vaccine opinion diffusion that addresses the dual propagation of positive and negative opinions, as investigated in the previous Chapter (see Section 3.2.1). For epidemic modeling, following Campbell and Salathé (2013), we utilize the SIR model for the disease spread stage, originally developed by Kermack and McKendrick (1927), which is widely used by researchers in epidemic modeling.

To begin, we derive the equations that govern the opinion dynamics, which are inherently probabilistic. The probability for an agent i to shift from a neutral state to an anti-vaccine state (Pr_i^-) or a pro-vaccine state (Pr_i^+) at time t is determined by three primary factors: (1) social network influence, where anti-vaccine neighbors (K_i^-) and pro-vaccine neighbors (K_i^+) each exert an influence with probability ω per time step; (2) environmental influence, such as media and campaigns, characterized by influence strengths μ_i^- and μ_i^+ for anti- and pro-vaccine information, respectively; and (3) the agent's historical exposure to both positive and negative information. Figure 4.1 presents a schematic depiction of the influence experienced by an agent i at a single time step.

We begin by analyzing the likelihood of an agent transitioning from a neutral state at time t to a negative opinion state at the next time step, $t+1$. As described in the previous chapter, for an agent to adopt a negative opinion, the agent needs θ more negative exposures than positive exposures, and vice versa for adopting a positive opinion. Therefore, the condition that an agent i needs to satisfy to transition to a negative opinion state is given by:

$$X_i \geq Y_i + \theta, \tag{4.1}$$

where X_i represents the total number of negative exposures, and Y_i represents the total number of positive exposures experienced by agent i. In addition, X_i consists of social negative exposures, denoted by X_i^ω, and negative general exposures, denoted by X_i^μ, represented as follows:

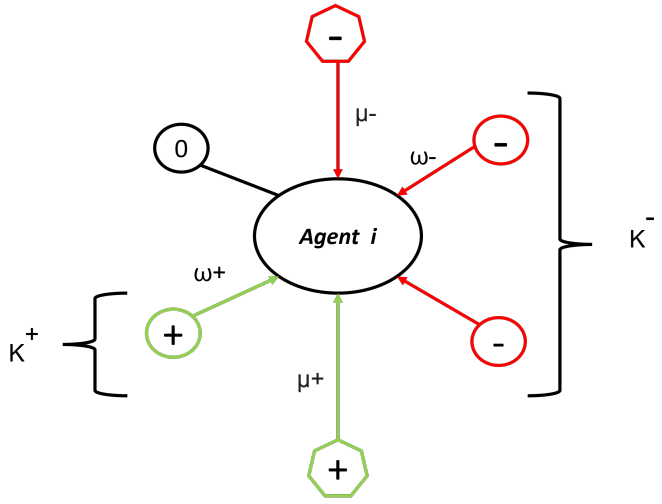

FIGURE 4.1: Schematic illustration of the influence exerted on an agent i, with polygons representing external exposures and circles representing social contacts for agent i. The $(-)$ sign denotes a negative influence exerted on node i, the $(+)$ sign represents a positive influence exerted on the node, and the 0 indicates neighbors with a neutral opinion. The probabilities for each influence are displayed over the links. K^- represents anti-vaccine neighbors, and K^+ represents pro-vaccine neighbors.

$$X_i = X_i^{\omega} + X_i^{\mu}. \tag{4.2}$$

The external negative exposure refers to influence originating from an external source (e.g. media, public messaging, etc.) and can influence an agent i with probability μ_i^-. We assume that this probability is uniform for all agents, such that $\mu_i^- = \mu^-$. Higher values of μ_i^-, with a maximum of 1, correspond to stronger influence on individuals, while lower values indicate weaker influence. In reality, the exposure may not be uniform and vulnerability to negative sentiment can depend on factors such as location, demographics, or background. These factors are not explicitly modeled in our framework and could be explored in future research with more detailed data. In contrast, social negative exposures refers to influence arising from social interactions. Specifically, it originates from the social contacts of an agent i who hold anti-vaccine opinions, each exerting influence with probability ω^-. A similar process applies to positive exposures Y_i, as follows:

$$Y_i = Y_i^{\omega} + Y_i^{\mu}, \tag{4.3}$$

where Y_i^{μ} represents the external positive exposure exerted on i with probability μ_i^+, corresponding to the influence exerted by intervention campaigns. Additionally, Y^{ω} denotes the social influence on agent i from social contacts who hold pro-vaccine opinions, each of whom can influence with probability ω^+.

The probability of an agent experiencing x negative social exposures, given a total of K_i^- negative neighbors, can be obtained using the Binomial distribution as follows:

$$p_i^-(X_i^\omega = x) = \begin{cases} \begin{pmatrix} K_i^- \\ x \end{pmatrix} (\omega^-)^x \cdot (1-\omega^-)^{K_i^- - x} & \text{if } 0 \le x \le K_i^- \\ 0 & \text{otherwise.} \end{cases} \tag{4.4}$$

Similarly, the probability of an agent experiencing y positive social exposures, given a total of K_i^+ positive neighbors, is obtained as follows:

$$p_i^+(Y_i^\omega = y) = \begin{cases} \begin{pmatrix} K_i^+ \\ y \end{pmatrix} (\omega^+)^y \cdot (1-\omega^+)^{K_i^+ - y}. & \text{if } 0 \le y \le K_i^+ \\ 0 & \text{otherwise.} \end{cases} \tag{4.5}$$

Furthermore, the probability of an agent experiencing negative general exposures can be obtained using the Bernoulli distribution, where an agent has probability μ_i^- of experiencing one negative exposure and probability $1-\mu_i^-$ of experiencing zero negative exposures. And similarly, for the probability of obtaining y positive general exposures. Consequently, the probability of obtaining a total of x negative exposures is given by:

$$P_i^-(X_i = x) = \mu_i^- \cdot p_i^-(x-1) + (1-\mu_i^-) \cdot p_i^-(x), \tag{4.6}$$

and the probability of obtaining a total of y positive exposures is as follows:

$$P_i^+(Y_i = y) = \mu_i^+ \cdot p_i^+(y-1) + (1-\mu_i^+) \cdot p_i^+(y). \tag{4.7}$$

To satisfy the condition in equation 4.1, where the number of negative exposures must exceed the number of positive exposures by θ, the overall negative probability can be determined by summing the probabilities of all occurrences that meet the condition as follows:

$$Pr_i^-(X_i \ge Y_i + \theta) = \sum_{a=0}^{K_i^- - \theta} \sum_{j=a+\theta}^{K_i^-} P_i^-(X_i = j) \cdot P_i^+(Y_i = a). \tag{4.8}$$

Up to this point, we have neglected any previous exposures that agent i may have encountered in prior time steps. The effective difference between positive and negative exposures is denoted by Δ_i and defined as $\Delta_i = \phi_i^+ - \phi_i^-$, where ϕ_i^+ and ϕ_i^- indicate the cumulative number of positive and negative exposures that agent i has

encountered up to time t, respectively. If agent i has already experienced a negative exposure but has not been exposed to any positive influence, then the agent becomes closer to adopting a negative opinion, requiring fewer additional negative exposures to transition. In other words, the threshold difference required to transition to the negative state in the next time step decreases. Conversely, if the agent has experienced a positive exposure without any negative exposure, the required threshold difference increases. This exposure information can be integrated by plugging in Δ_i into Eq. 4.8 as follows:

$$Pr_i^-(X_i \geq Y_i + \theta) = \sum_{a=0}^{\max(K_i^- - (\theta+\Delta_i),0)} \sum_{j=a+(\theta+\Delta_i)}^{\max(K_i^-,j)} P_i^-(X_i = j) \cdot P_i^+(Y_i = a), \tag{4.9}$$

We can identify a similar expression for Pr_i^+. Specifically, the probability that an agent adopts a positive opinion in the next time step requires the agent to have θ more positive exposures than negative exposures. Therefore, the condition that an agent i needs to satisfy to adopt a positive opinion is as follows:

$$Y_i \geq X_i + \theta \tag{4.10}$$

Consequently, the probability of an agent i converting to a positive state is as follows:

$$Pr_i^+(Y_i \geq X_i + \theta) = \sum_{a=0}^{\max(K_i^+ - (\theta-\Delta_i),0)} \sum_{j=a+(\theta-\Delta_i)}^{\max(K_i^+,j)} P_i^-(X_i = a) \cdot P_i^+(Y_i = j). \tag{4.11}$$

Given the state of the system at a particular time t, we derive two expressions to estimate: (1) the expected number of anti-vaccine opinion adopters at the next time step, based on Eq. 4.9; and (2) the expected number of pro-vaccine opinion adopters at the next time step, based on Eq. 4.11 . Predicting the long-term evolution of vaccine opinions is intricate due to the inherent stochasticity of these processes. Therefore, we develop one-step-ahead optimization strategies aimed at: (1) minimizing the expected number of individuals adopting anti-vaccine opinions in the near future (i.e., next time step); and (2) maximizing the expected number of individuals adopting pro-vaccine opinions in the next time step, based on the current data regarding vaccination attitudes within the social network.

Building on these foundations, we introduce our intervention campaigns to determine optimal resource allocation, as outlined in the following sections. Here, motivated by the continuous influence maximization technique presented in Romero Moreno et al. (2021), our goal is to optimally allocate the campaign budget, denoted as B, where $B = \mu^+ \times N$ and μ^+ represents the budget per individual, across the population over

time. In more detail, we aim to dynamically determine the optimal allocation μ_i^+ for each agent $i = 1, \ldots, N0$, where N0 represents the set of individuals holding a neutral opinion at time t, subject to a given campaign budget B.

The budget B reflects the available resources of the campaign and can be interpreted as the strength of campaign influence. A higher budget implies stronger influence, enabling more impactful interventions across the population. In practice, the budget can be interpreted as the number of individuals who can be positively influenced per unit of time. For example, this could be a health officer visiting households, speaking directly with residents and providing guidance. In this scenario, the officer can only reach a limited number of people per day. Alternatively, the campaign could involve sending emails or messages to individuals, a process that is much faster and allows more people to be contacted in the same period, although the resulting influence may be less pronounced.

To evaluate the performance of these optimization campaigns, we use two benchmark campaigns introduced in Chapter 3: the dynamic random campaign DynRandT and the best-performing heuristic DynAdvLocT campaign.

4.2.1 Minimizing Anti-vaccine Opinion Adopters

In this section, we introduce the optimization-based campaign aimed at minimizing the number of individuals who adopt an anti-vaccine opinion in the next time step. Let $\langle \Delta n^- \rangle$ represent the average number of anti-vaccine opinion adopters. Then, the expected number of anti-vaccine opinion adopters in the next time step, $\langle \Delta n^- \rangle_{(t+1)}$, based on the network state at time step t, can be obtained as follows:

$$\langle \Delta n^- \rangle_{(t+1)} = \sum_{i \in N^0} Pr_i^-, \tag{4.12}$$

Accordingly, this estimator is the objective function that we aim to minimize using optimization algorithms. In the next section, we assess the accuracy of the estimator by comparing its predictions with numerical simulation results. Additional details are provided within the section. To proceed with the optimization, let $\boldsymbol{\mu}^+$ represent a vector whose entries correspond to μ_i^+. Then, the budget-constraint optimization problem can be formalized with respect to $\boldsymbol{\mu}^+$ as follows:

$$\min_{\boldsymbol{\mu}^+} \langle \Delta n^- \rangle_{(t+1)}, \text{ s.t.} \qquad \sum_i \mu_i^+ \leq B, \quad 0 \leq \mu_i^+ \leq 1. \tag{4.13}$$

Given the linear nature of the objective function and the associated constraints, we leverage linear programming methods, particularly the simplex method (Press et al.,

2007), to derive the optimal allocation $\boldsymbol{\mu}^+$. Therefore, equation 4.12 can be further rewritten as:

$$\left\langle \Delta n^- \right\rangle = \mathbf{c}^\top \boldsymbol{\mu}^+ + c_0, \tag{4.14}$$

where $\mathbf{c}$ is defined as:

$$\mathbf{c} = \begin{bmatrix} \sum_{a=0}^{K_1^- - (\theta - \Delta_1)} \sum_{j=a+(\theta-\Delta_1)}^{K_1^-} P^-(j) \cdot \left(p^+(a-1) - p^+(a)\right) \\ \vdots \\ \sum_{a=0}^{K_{N^0}^- - (\theta - \Delta_{N^0})} \sum_{j=a+(\theta-\Delta_{N^0})}^{K_{N^0}^-} P^-(j) \cdot \left(p^+(a-1) - p^+(a)\right) \end{bmatrix}, \tag{4.15}$$

$\boldsymbol{\mu}^+ = \left[\mu_1^+, \mu_2^+, \ldots, \mu_{N^0}^+\right]^\top$, and $c_0 = \sum_{i \in N^0} \sum_{a=0}^{K_i^- - (\theta - \Delta_i)} \sum_{j=a+(\theta-\Delta_i)}^{K_i^-} P^-(j) \cdot p^+(a)$.

4.2.2 Maximizing Pro-vaccine Opinion Adopters

Unlike the goal in the previous section, which aimed to minimize the number of anti-vaccine opinion adopters, here we explore the impact of maximizing the number of pro-vaccine (vaccinated) individuals within the network on containing the negative influence, and the impact on the final epidemic size. Increasing vaccine coverage is a key objective for public health authorities seeking to achieve herd immunity, and this goal has attracted considerable attention in the literature exploring, for example, different immunization strategies (e.g., (Chen et al., 2022a)). The majority of these studies focus on the interaction between pathogen dynamics and vaccination with different vaccination policies (Yang et al., 2019b; Chen et al., 2022a) or incentive-based strategies (Zhang et al., 2014, 2017; Tatsukawa et al., 2021), with relatively limited consideration of the influence of vaccine misinformation and its associated negative behavioral responses (Prieto Curiel and González Ramírez, 2021), as well as the impact of interventions aimed at mitigating such misinformation. While increasing the number of vaccinated agents may appear to be an effective strategy, it remains unclear whether this approach effectively mitigates the spread of anti-vaccine influence and the emergence of large unprotected communities. As a result, it is uncertain whether it contributes more to reducing the final epidemic size compared to the minimization approach.

Thus, in this section, we introduce an optimization-based campaign aimed at maximizing the number of individuals who adopt a pro-vaccine opinion in the next time step. Let $\langle \Delta n^+ \rangle$ represent the average number of pro-vaccine opinion adopters, and N^0 represent the set of current agents with neutral opinion. Then, the expected

number of pro-vaccine opinion adopters in the next time step, $\langle \Delta n^+ \rangle_{t+1}$, can be obtained as follows:

$$\langle \Delta n^+ \rangle_{t+1} = \sum_{i \in N^0} Pr_i^+. \tag{4.16}$$

This represents our second objective function that we aim to maximize. To verify the accuracy of this estimator, we compare the estimated number of pro-vaccine opinion adopters with numerical results obtained through simulations in the next section. The budget-constrained optimization problem can then be formalized as follows:

$$\max_{\boldsymbol{\mu}^+} \langle \Delta n^+ \rangle_{(t+1)}, \text{ s.t.} \qquad \sum_i \mu_i^+ \leq B, \quad 0 \leq \mu_i^+ \leq 1. \tag{4.17}$$

Similarly to our approach above, given the linear nature of the objective function and the associated constraints, we leverage linear programming methods, particularly the simplex method (Press et al., 2007), to derive the optimal allocation $\boldsymbol{\mu}^+$. Therefore, equation 4.16 can be further rewritten as:

$$\left\langle \Delta n^+ \right\rangle = \mathbf{c}^\top \boldsymbol{\mu}^+ + c_0, \tag{4.18}$$

where $\mathbf{c}$ is defined as:

$$\mathbf{c} = \begin{bmatrix} \sum_{a=0}^{K_1^+ - (\theta - \Delta_1)} \sum_{j=a+(\theta-\Delta_1)}^{K_1^+} P^-(a) \cdot \left(p^+(j-1) - p^+(j)\right) \\ \vdots \\ \sum_{a=0}^{K_{N^0}^+ - (\theta - \Delta_{N^0})} \sum_{j=a+(\theta-\Delta_{N^0})}^{K_{N^0}^+} P^-(a) \cdot \left(p^+(j-1) - p^+(j)\right) \end{bmatrix}, \tag{4.19}$$

and $c_0 = \sum_{i \in N^0} \sum_{a=0}^{K_i^+ - (\theta - \Delta_i)} \sum_{j=a+(\theta-\Delta_i)}^{K_i^+} P^-(a) \cdot p^+(j)$.

4.2.3 Numerical Verification

To verify the accuracy of the estimators in Eqs. 4.12 and 4.16, we compare the estimated number of anti-vaccine opinion adopters (Eq. 4.12) and pro-vaccine opinion adopters (Eq. 4.16) with numerical results obtained through simulations. To this end, we run the simulator and capture the network state at distinct time steps t during the opinion diffusion phase. We then analytically compute the expected number of anti-vaccine adopters, $\langle \Delta n^- \rangle$, for the next time step using Eq. 4.12, and the expected number of pro-vaccine adopters, $\langle \Delta n^+ \rangle$, for the next time step using Eq. 4.16. To obtain the numerical value of anti-vaccine adopters at time $(t+1)$, we run the

simulator multiple times for this specific time step and compute the average across these simulations. We then compare the numerical result with the analytical estimate to assess the accuracy of the estimators.

Figure 4.2 illustrates a comparison between the analytical results obtained from Eq. 4.12 and the numerical results obtained from the simulations for anti-vaccine opinion adopters. Similarly, Figure 4.3 illustrates a comparison between the analytical results obtained from Eq. 4.16 and the numerical results obtained from the simulations for pro-vaccine opinion adopters. Here, a uniform influence allocation is assumed for all agents, with $\mu_i^+ = \mu_i^- = 0.001$ and $\omega^{+,-} = 0.006$. As shown in the figures, the analytical and numerical numbers of both anti-vaccine and pro-vaccine opinion adopters are identical within error estimates.

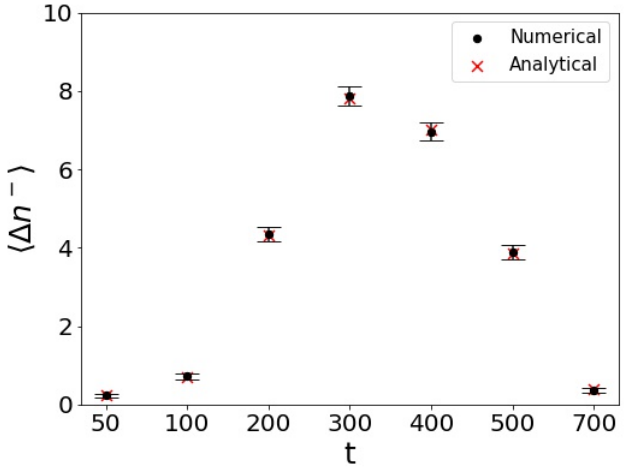

FIGURE 4.2: A comparison between the analytical and numerical values of the number of anti-vaccine adopters, $\langle \Delta n^- \rangle$, obtained using Equation 4.12 (in red) and simulations (in black), respectively. Each date point corresponds to the $\langle \Delta n^- \rangle$ at time $(t+1)$ based on the network state at time t during the opinion diffusion phase. The numerical values are averaged over 500 realizations, with error bars corresponding to the 95% confidence interval.

4.3 Results

In this section, we analyze the control performance of the proposed optimization strategies. We begin by presenting the results of the final epidemic size in comparison to the benchmark campaigns, along with the impact on the final number of pro- and anti-vaccine opinion adopters (Sections 4.3.1 to 4.3.3). Next, we extend the analysis to a more realistic setting where early intervention is not feasible, assessing the effectiveness of the campaigns under delayed response scenarios (Section 4.3.4). Finally, Section 4.3.5 provides a structural analysis of anti-vaccine communities and examines their influence on mitigation outcomes.

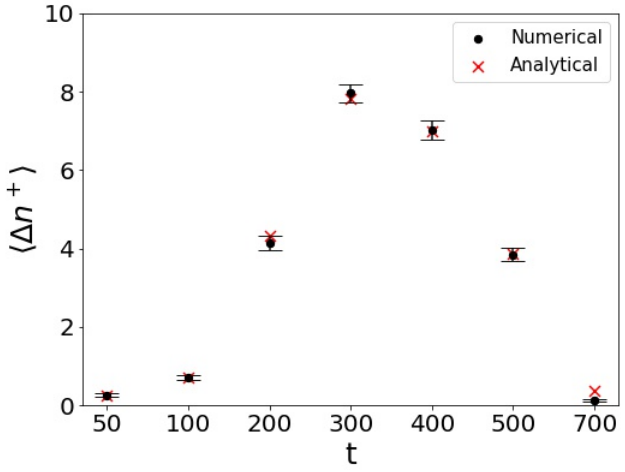

FIGURE 4.3: A comparison between the analytical and numerical values of the number of pro-vaccine adopters, $\langle \Delta n^+ \rangle$, obtained using Equation 4.16 (in red) and simulations (in black), respectively. Each date point corresponds to the $\langle \Delta n^- \rangle$ at time $(t+1)$ based on the network state at time t during the opinion diffusion phase. The numerical value is averaged over 500 realizations. The numerical values are averaged over 500 realizations, with error bars corresponding to the 95% confidence interval.

We first present the results of the system dynamics in a simplified setting. We consider an intervention campaign launched simultaneously with the onset of negative influence, considering an equal-budget scenario where the strengths of both negative and positive influences are equal, i.e. $\mu^- = \mu^+$. However, this assumption does not accurately reflect real-world scenarios, where interventions typically occur in response to the ongoing spread of vaccine misinformation, such as in the cases of the measles vaccine (Hotez et al., 2020) and the Covid-19 vaccine (Pierri et al., 2022). Therefore, to capture a more realistic setting, we next consider a scenario involving late interventions, in which strategies are implemented in the presence of existing negative influence propagation. In the late intervention scenario, the starting time differs between the negative general exposure and intervention campaigns, with exposure to negative influence beginning earlier than the intervention campaigns. We denote the starting times of the negative exposure and intervention campaigns as t_s^- and t_s^+, respectively.

For all experiments, we maintain the same network settings and diffusion parameters as in Chapter 3 for comparative analysis. We follow a similar configuration to that used in (Campbell and Salathé, 2013), considering a small-world network (Watts and Strogatz, 1998) with size $N = 5000$, an average degree $\langle k \rangle = 10$, and a rewiring probability $p = 0.01$. We use an opinion formation threshold of $\theta = 2$ as we consider complex contagion in which agents require more than one exposure to adopt an opinion. For the SIR model parameters, as in Campbell and Salathé (2013), we set the infection rate $\beta = 0.1$, recovery rate $\gamma = 0.1$, and seed set $I_0 = 1$. In all figures presented, the error bars and shaded area correspond to the 95% confidence intervals.

We use the random DynRandT campaign and the best-performing DynAdvLocT campaign introduced in Chapter 3 as benchmarks. Furthermore, for all optimization campaigns if the budget B is not fully utilized by the campaign, we randomly choose T nodes and distribute the remaining budget equally among them. Unless otherwise stated, the results correspond to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega^{+,-} = 0.006$, and $T = 5$. The results for each scenario are averaged over 50 realizations, each of which involves generating a network and simulating the opinion diffusion process followed by the disease spread.

4.3.1 Minimizing Anti-vaccine Opinion Adopters

To assess the control performance of the campaign aimed at minimizing anti-vaccine opinion adopters, referred to as Min_FK, we present the evolution of the number of anti-vaccine $\langle n^- \rangle$ and pro-vaccine $\langle n^+ \rangle$ opinion adopters, along with the corresponding epidemic size. Figure 4.4 presents the results of this campaign compared to the benchmark campaigns, where (A) shows the evolution of the number of anti-vaccine opinion adopters, (B) shows the evolution of the number of pro-vaccine opinion adopters, and (C) shows the corresponding epidemic size.

As depicted in Fig. 4.4(A), the Min_FK campaign, represented by the green color, results in a slower diffusion of negative influence for an extended period. However, it ultimately fails to contain the overall spread of this influence in the network, leading to a substantial increase in anti-vaccine opinion adopters compared to the benchmark campaigns and, as a result, a larger epidemic size as shown by the green bar in Fig. 4.4(C). Moreover, as can be seen in Fig. 4.4(B), the number of vaccinated individuals influenced by this campaign is very low, highlighting the algorithm's limitation in promoting vaccination.

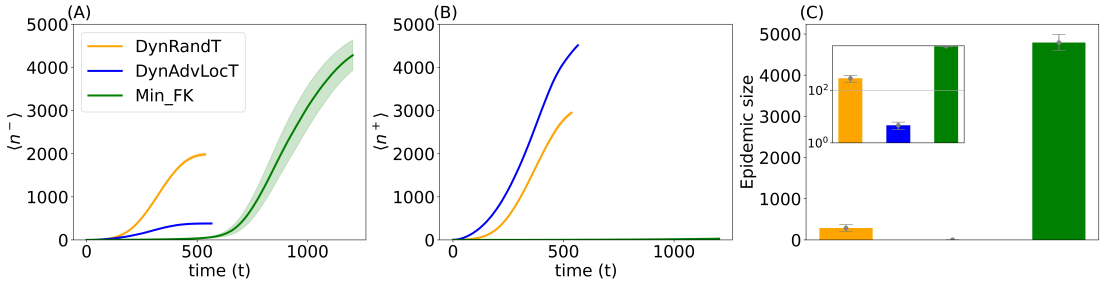

FIGURE 4.4: Final epidemic size and opinion evolution for the minimizing anti-vaccine opinion adopters campaign starting at time $t_s^+ = 0$ compared to the DynRandT and DynAdvLocT campaigns. Panel (A) illustrates the evolution of anti-vaccine opinion adopters, panel (B) shows the evolution of pro-vaccine opinion adopters (vaccinated), and panel (C) illustrates the resulting epidemic size. The results are related to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

To understand the poor performance of the Min_FK campaign, we analyze which agents are targeted by this campaign. For this purpose, we capture the neighborhood

structure of candidate agents to determine who received campaign allocation. Figure 4.5 provides snapshots of the network at different time steps during the opinion propagation phase. Each panel depicts the neighborhood structure of the uncommitted neutral agents, where the x-axis represents the number of anti-vaccine neighbors, the y-axis represents the number of pro-vaccine neighbors, and the cells represent the average allocation for agents in the corresponding neighborhood configuration. The top panels (Figs. 4.5(A)-(D)) correspond to the Min_FK campaign investigated in this section [1]. As seen from the average allocation in this figure, the campaign focuses on nodes with the highest number of anti-vaccine neighbors, represented by non-purple cells. For example, at $t = 500$ in Fig.4.5(C), the allocation is directed to nodes with one anti-vaccine neighbor, and at $t = 1100$ in Fig.4.5(D), to those with ten anti-vaccine neighbors. In all cases, the targets are the agents with the highest number of adjacent negative neighbors. Additionally, we observe that it further targets agents influenced by the opponent— those with negative exposure (i.e., $\Delta = -1$, the effective difference between positive and negative exposures, $\Delta_i = \phi_i^+ - \phi_i^-$)—as they are most likely to adopt a negative opinion in the next time step. For example, see the Fig.4.5(D) at $t = 1100$, where it prioritizes agents with ten anti-vaccine neighbors over those with more.

By targeting agents with a greater number of anti-vaccine neighbors, the campaign effectively protects the most susceptible agents from negative influence, stopping them from shifting to an anti-vaccine stance. Nevertheless, the failure of the Min_FK campaign to contain the overall negative diffusion over time stems from two factors: the campaign targeting mechanism and the limited budget of the positive campaign. As the targeting criteria of the Min_FK campaign are based not only on neighborhood structure but also on previous exposures, the campaign prioritizes individual protection over promoting positive influence. It avoids individuals who have been positively targeted before, unless they are at risk of converting to the negative state again. As a result, these agents remain neutral and susceptible to negative influence, and the campaign fails to provide sufficient exposures to positively influence them to get vaccinated. Recall that, due to complex contagion, agents require multiple exposures to be persuaded. Consequently, since prior knowledge limits the campaign's efficacy by providing inadequate exposure to the population, it also fails to seed positive influence in the network, thereby hindering its ability to combat the diffusion of negative influence. Additionally, the negative influence is amplified by an external negative exposure, which continues to impact the population randomly.

[1] The bottom panels correspond to the Min_NK campaign, another setting that will be discussed later in Section 4.3.3. The two analyses are presented in the same figure for comparison purposes.

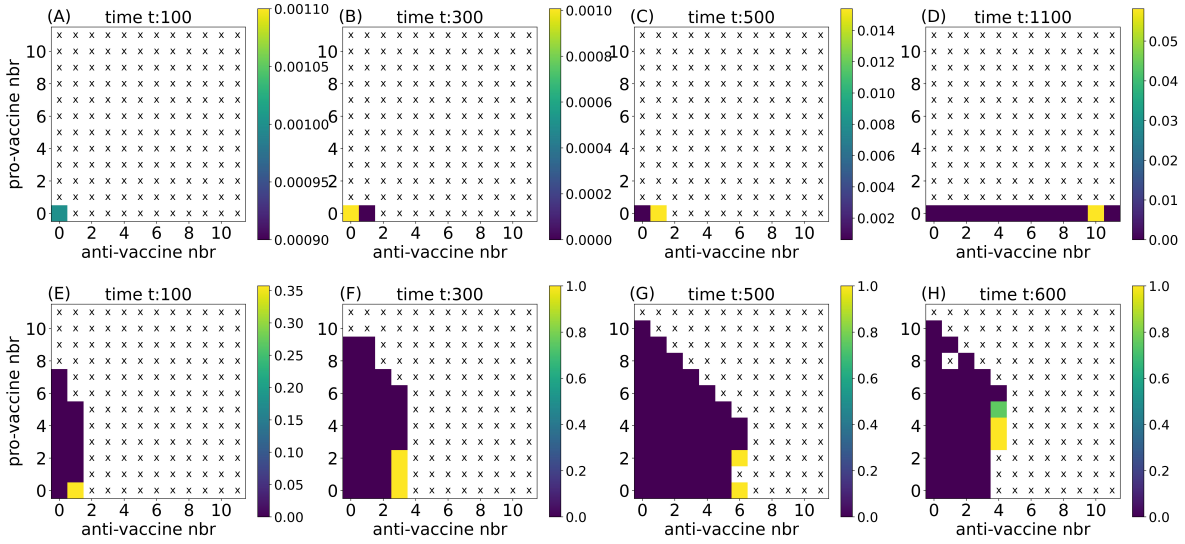

FIGURE 4.5: The targeting scheme for the minimizing anti-vaccine opinion adopters campaign is illustrated with complete knowledge (top panel) and incomplete knowledge (bottom panel). These panels illustrate the neighborhood structure at different time steps t during the opinion propagation phase, displaying the number of pro-vaccine and anti-vaccine neighbors for uncommitted neutral agents. Each cell represents the average allocation for agents with x anti-vaccine neighbors and y pro-vaccine neighbors. An 'x' in the figure further denotes that no agent was found in this particular neighborhood configuration. Targeted agents are those with a non-zero allocation (indicated by non-purple colors).

4.3.2 Maximizing Pro-vaccine Opinion Adopters

In this section, we examine the impact of maximizing the number of pro-vaccine opinion adopters (Max_FK), representing vaccinated individuals within the network. To evaluate the effectiveness of this campaign, we present the results in Fig. 4.6, where Fig. 4.6(A) shows the evolution of anti-vaccine opinion adopters, Fig. 4.6(B) shows the evolution of pro-vaccine opinion adopters, and Fig.4.6(C) presents the corresponding epidemic size. As shown in Fig.4.6 (C), although this campaign performs better than Min_FK in Fig. 4.4, it fails to effectively reduce the final epidemic size, with the benchmark campaigns significantly outperforming it and achieving a greater reduction.

To understand the observed control performance of this campaign, it is essential to analyze its targeting strategy. Following the approach used in the previous section, we examine the neighborhood structure of candidate agents to identify those who received influence allocation. Figure 4.7 presents snapshots of the network at different time steps during the opinion propagation phase. Each panel illustrates the neighborhood configuration of candidate agents (neutrals), where the x-axis

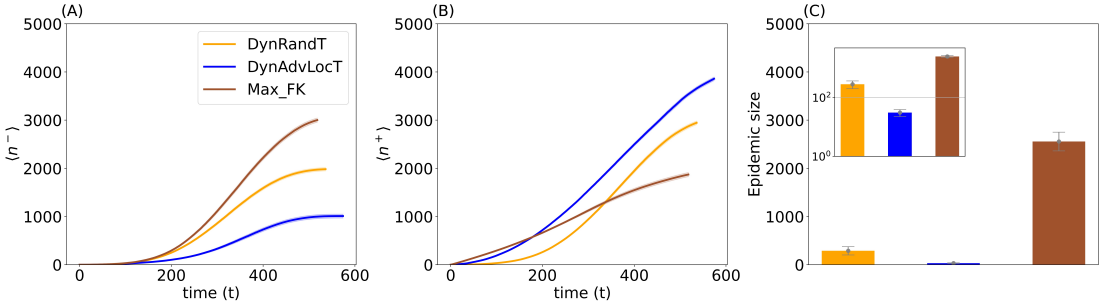

FIGURE 4.6: Final epidemic size and opinion evolution for the maximizing pro-vaccine opinion adopters campaign starting at time $t_s^+ = 0$ compared to the DynAdvLocT and DynRandT campaigns. Panel (A) illustrates the evolution of anti-vaccine opinion adopters, panel (B) illustrates the evolution of pro-vaccine opinion adopters (vaccinated), and panel (C) illustrates the resulting epidemic size. The results are related to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

represents the number of anti-vaccine neighbors, the y-axis represents the number of pro-vaccine neighbors, and the cells indicate the average allocation for agents in the corresponding neighborhood configuration. The top panel (Figs.4.7(A)-(D)) corresponds to the Max_FK campaign investigated in this section[2].

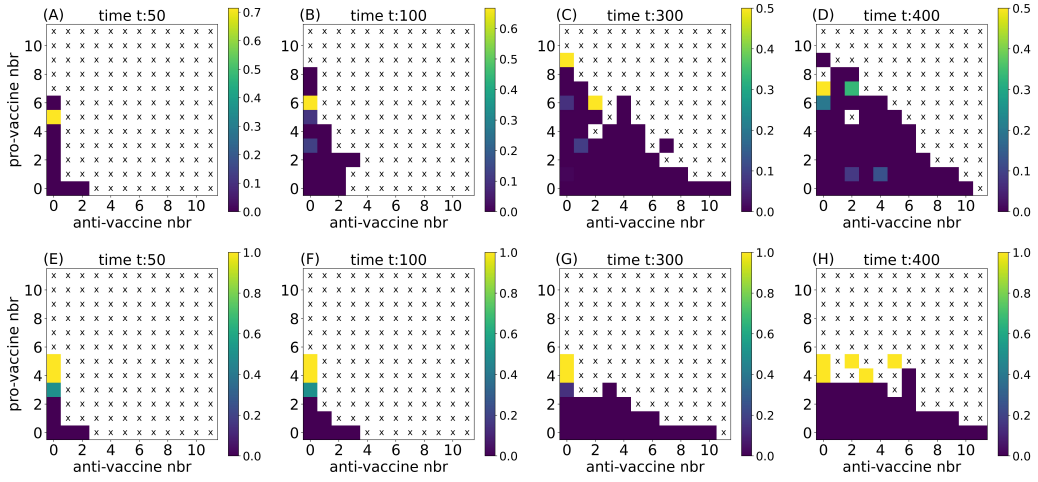

FIGURE 4.7: The targeting scheme for the maximizing pro-vaccine opinion adopters campaign is illustrated with complete knowledge (top panel) and incomplete knowledge (bottom panel). These panels illustrate the neighborhood structure at different time steps t during the opinion propagation phase, displaying the number of pro-vaccine and anti-vaccine neighbors for uncommitted neutral agents. Each cell represents the average allocation for agents with x anti-vaccine neighbors and y pro-vaccine neighbors. An 'x' in the figure further denotes that no agent was found in this particular neighborhood configuration. Targeted agents are those with a non-zero allocation (indicated by non-purple colors).

[2]The bottom panels correspond to the Max_NK campaign, another setting that will be discussed later in Section 4.3.3. The two analyses are presented in the same figure for comparison purposes.

We observe that the campaign primarily targets agents who are adjacent to pro-vaccine neighbors, as indicated by the non-purple cells across all panels (Figs.4.7(A)-(D)). Among these, it prioritizes those with the highest number of pro-vaccine neighbors and the fewest anti-vaccine neighbors. The campaign further targets agents with $\Delta = 1$ the effective difference between positive and negative exposures, $\Delta_i = \phi_i^+ - \phi_i^-$), representing individuals who have been exposed to positive information and are therefore most likely to adopt the pro-vaccine opinion in the next time step. Since the algorithm focuses on agents most likely to convert to the positive state, it generally avoids targeting individuals adjacent to anti-vaccine influence, as they are less likely to respond positively and adopt the desired opinion.

As a result, this algorithm fosters the spread of positive influence and facilitates the growth of pro-vaccine communities. However, by not prioritizing targets in close proximity to negatively influenced areas of the network, it limits its effectiveness in countering the spread of anti-vaccine influence in high-risk regions. This behavior is evident in Figs.4.7(A) and (B), where the selected targets are primarily agents with no anti-vaccine neighbors, indicating that the campaign focuses on areas distant from zones of negative diffusion. In the later stages, as both negative and positive influence spread more widely, the algorithm may begin to target individuals in closer proximity to negative influence who most closely meet the selection criteria. However, this occurs too late, as large anti-vaccine communities may have already formed, making it difficult to contain the negative propagation. For example, see Figs.4.7 (C) and (D) where agents are targeted near anti-vaccine neighbors. These figures also show the presence of agents with a high number of anti-vaccine neighbors (x-axis), indicating clustering among anti-vaccine agents.

Targeting agents with pre-existing pro-vaccine neighbors is not an efficient strategy. This approach neither helps initiate positive contagion in new areas of the network nor stops the spread of negative influence in regions already affected. This is similar to what we observed in the previous campaign Min_FK, where targeting agents with the highest number of anti-vaccine neighbors offered protection to those individuals but failed to contain the broader propagation of negative influence. Therefore, this approach is inefficient, as it does not fulfill our primary goal of preventing the clustering of anti-vaccine opinion adopters. Our findings show that anti-vaccine individuals are distributed across 15 distinct communities by Max_FK campaign. While this represents an improvement over the Min_FK campaign, which produces a higher number of anti-vaccine individuals concentrated in a single large cluster, the DynAdvLocT campaign performs even better, resulting in 164 smaller communities [3]. This suggests that increasing vaccine coverage alone may not guarantee the prevention of outbreaks, especially when unvaccinated individuals form large

[3]Further analysis of the impact of different campaigns on anti-vaccine communities will be provided in Section 4.3.5.

communities, as this facilitates the widespread transmission of infectious diseases within these communities. This aligns with previous studies (Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Gromis and Liu, 2022) who highlight that the clustering of anti-vaccine individuals poses a significant risk and can undermine efforts to achieve herd immunity. This is also reflected in ongoing measles outbreaks which continue to occur globally across different countries (Do and Mulholland, 2025). For instance, the United States reported 49 measles outbreaks by December 2025, even though vaccination coverage remained high at 92% in 2023–2024, down from 95% in 2019–2020 (CDC, 2025).

4.3.3 What If We Do Not Consider Previous Exposures?

In previous sections, we find that the targeting schemes of both campaigns depend not only on agents' neighborhood structures but also primarily on their prior exposures. The minimization strategy focused on protecting those likely to be influenced negatively, while the maximization strategy targeted agents most likely to be persuaded. Consequently, providing complete knowledge of agents to the campaigns resulted in neither minimizing anti-vaccine opinion adopters nor maximizing pro-vaccine opinion adopters being efficient strategies to counteract the spread of negative influence. This ultimately led to the formation of a large community of unprotected individuals and as a result a high epidemic size. This raises the question: what if prior exposures are not considered, and only the current state of the network is provided to the algorithms? In such a case, the targeting criteria would be driven solely by the agents' neighborhood structures.

Therefore, in this section, we disregard prior knowledge of previous exposures and assume that historical exposure information is unknown. We aim to explore the effectiveness of the campaign under this incomplete knowledge setting. We model incomplete knowledge by setting $\Delta_i = 0$ in the objective functions in Eq. 4.12 and Eq.4.16. The suffixes NK and FK are used to denote campaigns with incomplete and full knowledge settings, respectively.

To evaluate the effectiveness of this approach, we present in Fig. 4.8 the results for both campaigns with full knowledge and incomplete knowledge settings compared to the benchmark campaigns DynRandT and DynAdvLocT. Figure 4.8(A) shows the evolution of anti-vaccine opinion adopters, Fig. 4.8(B) shows the evolution of pro-vaccine opinion adopters, and Fig.4.8(C) presents the corresponding epidemic size.

As shown in Fig.4.8(C), the Min_NK campaign (black bar) achieves significantly better performance when neglecting historical exposures, resulting in a substantial reduction in epidemic size. This campaign minimizes the number of anti-vaccine agents and

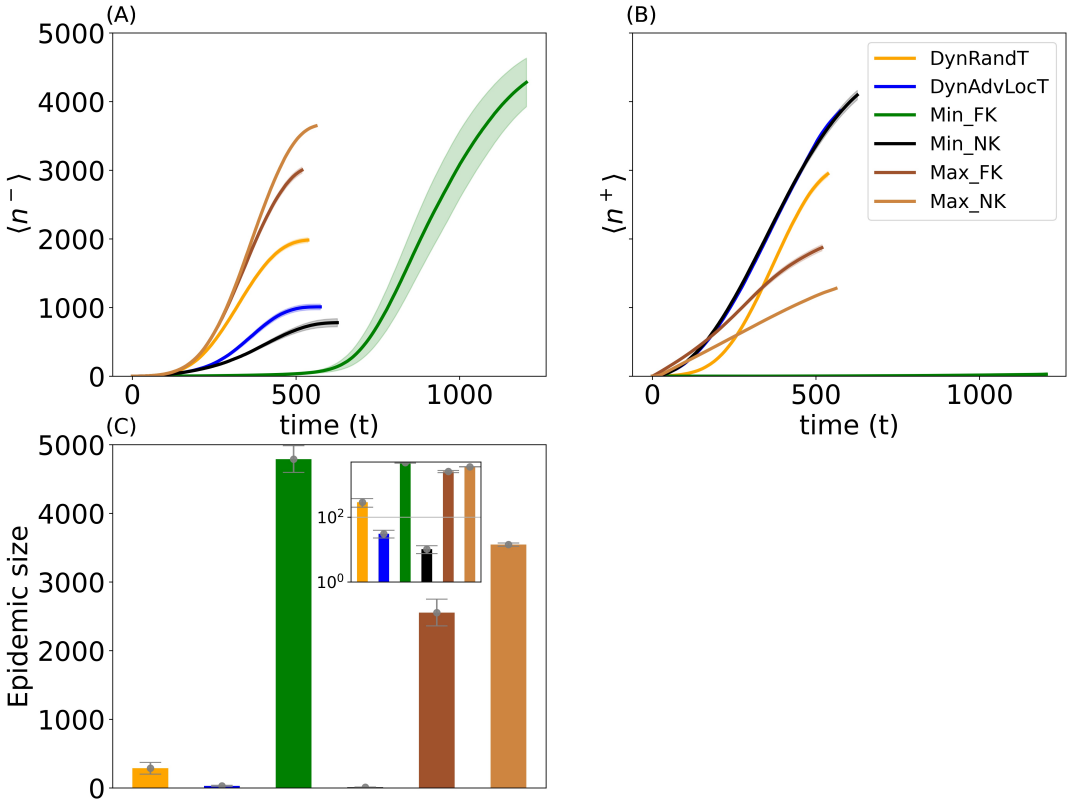

FIGURE 4.8: Final epidemic size and opinion evolution for minimization and maximization campaigns starting at time $t_s^+ = 0$. Results are shown for full knowledge (FK) and incomplete knowledge (NK) settings compared to the DynRandT and DynAdvLocT campaigns. Panel (A) illustrates the evolution of anti-vaccine opinion adopters, panel (B) illustrates the evolution of pro-vaccine opinion adopters (vaccinated), and panel (C) illustrates the corresponding epidemic size. The inset plot in panel (C) presents the same results as panel (C) on a log-scaled axis to illustrate variations. The results are related to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

bolsters the propagation of positive influence, with a marked increase in vaccinated individuals (see black curves in Figs.4.8(A) and (B)). In this campaign, resource allocation depends solely on the neighborhood structure, regardless of an agent's previous exposures. As a result, agents can be targeted multiple times, and the repeated exposures delivered to those susceptible to negative influence are sufficient to persuade them to vaccinate, thereby facilitating the initiation of positive contagion within the network. This contrasts with the previous setting, (Min_FK), which avoided pre-targeted agents and was therefore less effective at promoting vaccine uptake (see green lines and bar of Fig.4.8). The Min_FK campaign performance is comparable to the DynAdvLocT campaign, represented by blue color, with both resulting in a

significantly smaller epidemic size.

The targeting scheme (see Figs. 4.5(E)-(H)) shows that the campaign with incomplete knowledge setting still prioritizes nodes with a high number of anti-vaccine neighbors. For example, at $t = 300$ in Fig.4.5(F), the campaign targets nodes with three anti-vaccine neighbors, and at $t = 600$ in Fig.4.5(H), those with four anti-vaccine neighbors. In all cases, the targets are the agents with the highest number of adjacent negative neighbors. However, the network configuration reveals growth in pro-vaccine neighbors (y-axis), see Figs.4.5(E)-(H), indicating the campaign's success in initiating positive contagion in proximity to negative clusters, in contrast to the complete knowledge setting in the top panels (Figs. 4.5(A)-(D)). This dual effect—protection and promotion—helps contain the negative spread and initiate positive contagion early, see for example the bottom panels with ($t < 500$) in Figs.4.5(E) and (F). As a result, this campaign effectively disrupts the connectivity between anti-vaccine communities and ultimately leads to a significant reduction in epidemic size, as illustrated by the black bar in Fig.4.8(C). Therefore, in the rest of this study, we consider the incomplete knowledge setting for this campaign due to its superior performance.

For maximizing campaigns, as shown in light and dark brown bars in Fig.4.8(C), this approach fails to effectively minimize the final epidemic size, even without knowledge of historical exposures. In fact, incorporating exposure history in this approach (Max_FK campaign) results in a slightly higher number of pro-vaccine opinion adopters and, consequently, a slight reduction in the epidemic size (see dark brown in Fig.4.8). It also provides better control than the Min_FK campaign, leading to a smaller epidemic size (see the dark brown and green bars in Fig.4.8(C)). Nevertheless, the Min_NK campaign substantially outperforms all of them.

By analyzing the targeting scheme of this campaign in Figs.4.7(E)-(H), we observe that it continues to target agents adjacent to pro-vaccine neighbors, as indicated by the non-purple cells across all panels. As demonstrated in Section 4.3.2, the campaign prioritizes individuals with the highest number of pro-vaccine neighbors and the fewest anti-vaccine neighbors. This implies that the campaign primarily operates near existing positive clusters and focuses on their expansion. However, in the incomplete knowledge setting, it no longer prioritizes those with $\Delta = 1$, which represents individuals who have previously been exposed to positive information. Including prior exposure information helps to slightly broaden the campaign's reach by distributing attention across multiple regions in the network, even though these are still in proximity to positive clusters. In contrast, removing this knowledge reduces that effect, causing the campaign to concentrate on fewer areas, which ultimately leads to poorer performance.

Our findings suggest that minimizing the number of anti-vaccine opinion adopters leads to more effective containment of negative diffusion than strategies aimed at maximizing pro-vaccine adoption. This approach is able to promote vaccination near areas of negative influence, helping to mitigate the connectivity and growth of emerging anti-vaccine communities and resulting in a significant reduction in epidemic size. Although the anti-vaccine minimization algorithm initially exhibits poor containment with full knowledge settings, its effectiveness is substantially improved by disregarding prior exposure history. This enhancement enables the algorithm to perform more effectively by delivering sufficient repeated exposures to the population, thereby positively influencing individuals and initiating positive contagion within the network. In contrast, the maximizing approach operates primarily near existing pro-vaccine regions to amplify their influence, often distant from negative regions. As a result, anti-vaccine communities continue to grow, ultimately leading to a larger epidemic size.

4.3.4 Impact of Intervention Timing

So far, we have assumed that the intervention takes place in a neutral network, where all individuals initially hold a neutral opinion, and that the campaign is launched concurrently with the onset of negative opinion diffusion. This setup allows us to gain insights into the effectiveness of different campaigns in controlling the spread of anti-vaccine attitudes and the resulting epidemic size. However, in reality, interventions often take place in response to the ongoing spread of rumors and misinformation that has already influenced a portion of the population. For instance, misinformation regarding the measles vaccine has led to a decline in vaccination coverage, resulting in a resurgence of the disease (Hotez et al., 2020), and highlighting the need for tailored intervention campaigns (Thompson et al., 2023). Similarly, misinformation has negatively affected the acceptance of the Covid-19 vaccine (Pierri et al., 2022).

Therefore, in this section we relax the assumption of launching the campaign in a neutral network at $t_s^+ = 0$ and explore a more realistic scenario of late intervention, where the intervention is started in a network already affected by negative influence. Can we contain this spread and mitigate the growth of the anti-vaccine communities? What will the shape of the campaigns be in mitigating the spread of anti-vaccine influence? In this scenario, we let the negative general exposure start influencing the network at $t_s^- = 0$, while our intervention campaigns start later at time $t_s^+ = t$. At each starting time t_s^+, the level of negative dominance varies, represented by the number of existing anti-vaccine opinion adopters, which reflects different proportions of the population influenced by negative information. We aim to evaluate the effectiveness of our intervention campaigns under varying degrees of negative dominance, with an

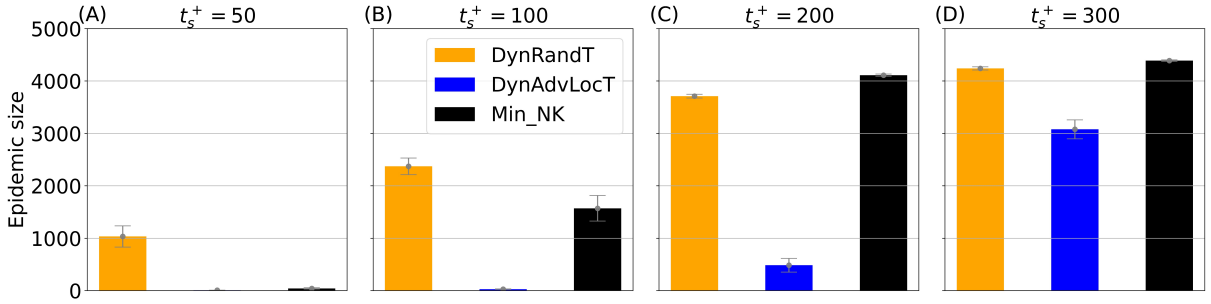

FIGURE 4.9: Final epidemic size obtained by applying the Min_NK campaign at time t_s^+ compared to benchmark campaigns DynAdvLocT and DynRandT. The intervention starts at (A) $t_s^+ = 50$, (B) $t_s^+ = 100$, (C) $t_s^+ = 200$, and (D) $t_s^+ = 300$. The results are related to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

equal-budget scenario in which both negative and positive campaigns operate with equal budget. We will explore the impact of campaign budget on control effectiveness in Chapter 6.

To evaluate the control efficacy of the campaigns, Fig. 4.9 presents the final epidemic size results for interventions starting at different time steps (A) $t_s^+ = 50$, (B) $t_s^+ = 100$, (C) $t_s^+ = 200$, and (D) $t_s^+ = 300$. These correspond to approximately 0.1%, 1%, 10%, and 26% of the population, respectively, having adopted the anti-vaccine opinion at the time of intervention. We consider the best-performing campaigns, DynAdvLocT and Min_NK, in addition to DynRandT as a baseline campaign. Since the maximization campaigns demonstrated poor performance, we excluded this approach from the current analysis and all subsequent evaluations. As shown in Fig. 4.9(A), when we intervene at time $t_s^+ = 50$, with only a few individuals having adopted the anti-vaccine opinion, the campaigns remain effective and achieve a significant reduction in epidemic size. While all campaigns, including DynRandT, provide relatively notable control, the best-performing strategies are the DynAdvLocT and Min_NK campaigns. However, control performance begins to degrade the longer the intervention is delayed. We find that at $t_s^+ = 100$ in Fig. 4.9(B), with a higher number of existing anti-vaccine opinion adopters, the Min_NK campaign results in a larger epidemic size compared to $t_s^+ = 50$ scenario in Fig. 4.9(A). Nevertheless, it still performs well, preventing infection in up to approximately 70% of the population. The DynAdvLocT maintains even more effective control in this scenario compared to the Min_NK campaign. In Fig. 4.9(C) with the intervention starting at $t_s^+ = 200$, control performance degrades significantly for both Min_NK and DynRandT, with the latter performing slightly better. The DynAdvLocT campaign shows superior performance up to this point, but its control also begins to degrade, although it still provides strong control overall. In Fig. 4.9(D) with intervention at $t_s^+ = 300$, the effectiveness deteriorates further, with the best outcome achieved by the DynAdvLocT campaign, where approximately 62% of the population becomes infected.

A delayed intervention leads to an earlier onset of broader negative social contagion and increased clustering among anti-vaccine opinion adopters. These clusters grow steadily the longer the intervention is delayed, resulting in the formation of large anti-vaccine communities. We will further analyze the distribution of anti-vaccine opinion adopters and the structure of anti-vaccine communities in the following section. Once these anti-vaccine clusters have developed, also a pattern observed among unvaccinated populations (Lieu et al., 2015), it becomes challenging to curb their expansion and mitigate their adverse effects with a limited budget. This consequently decreases the efficacy of the intervention campaigns in containing the growth of anti-vaccine communities and reducing the epidemic size. In addition, we observe that campaigns' resilience to the level of existing negative dominance differs, with DynAdvLocT demonstrating greater resilience than Min_NK, whose performance declines earlier.

We find that the earlier the intervention is applied, the more effectively negative diffusion can be countered. Early intervention enables the campaigns to initiate positive contagion within the network. In such scenarios, the social network configuration of the targeted individuals facilitates the spread of their positive influence within their communities. In other words, agents at early stages may have some anti-vaccine neighbors in their social network, but they also have neutral individuals who have not yet formed a vaccine stance and can potentially be positively influenced by the targeted agent. In contrast, an intervention implemented at a very late stage impedes the campaign's targeting efforts, as targeted agents are more likely to be encircled by numerous negative exposures within their social networks, making the negative influence far stronger than the campaign's impact. Moreover, even in situations where agents are positively influenced, disseminating their influence within a community already saturated with negative influence becomes significantly more challenging. This is also consistent with the findings from Chapter 3, specifically the results of DynLocT and DynAdvLocT in Section 3.4.4, which show that targeting agents with a larger number of anti-vaccine neighbors (higher values of ζ) leads to poorer performance compared to targeting agents with smaller values of ζ. This indicates the inherent limitations of the campaigns and highlights the necessity for further improvements.

4.3.5 Analysis of Anti-vaccine Communities

In this section, we examine the structure and evolution of anti-vaccine communities in two distinct states. First, we explore the clustering patterns of anti-vaccine opinion adopters prior to the application of intervention campaigns, in order to understand the reduced effectiveness observed across these strategies with late interventions. Second, we investigate the impact of different campaign strategies on the resulting population

composition, and how each strategy contributes to mitigating the connectivity among anti-vaccine communities. Our goal is to understand how different targeted interventions influence the disruption of large, connected anti-vaccine clusters.

As shown in the previous section, the longer the intervention is delayed, the more challenging the situation becomes. We hypothesize that this decline in performance is due to the emergence of large anti-vaccine communities that continue to grow and become increasingly difficult to control. Therefore, we perform an analysis of existing anti-vaccine communities prior to the initiation of the intervention. Fig. 4.10 presents the structure of these communities at different intervention times, where Fig. 4.10 (A) shows the average size of existing anti-vaccine communities, while Fig. 4.10 (B) shows the average number of these communities.

As shown in Fig. 4.10 (A), the longer the intervention is delayed, the larger these communities become. At $t_s^+ = 50$, represented by the brown bars, both the number and size of anti-vaccine clusters are relatively small. This explains why our campaigns, i.e., Min_NK and DynAdvLocT, remain highly effective at this point, resulting in a small epidemic size (see Fig. 4.9(A)). At $t_s^+ = 100$ (olive bars) in Fig. 4.10 (A), although the average size of anti-vaccine communities remains relatively small, more clusters have emerged in the network (see Fig. 4.10(B)). This proliferation of communities makes containment more challenging, particularly for the Min_NK campaign, as observed in Fig. 4.9(B). At $t_s^+ = 200$ (orange bars) in Fig. 4.10(A), with nearly 10% of the population having adopted anti-vaccine opinions, we observe that these individuals begin forming larger communities, and their number also increases significantly (as shown in Fig. 4.10(B)). Consequently, this significantly reduces the effectiveness of Min_NK campaign and results in a high epidemic size (see Fig. 4.9(C)). By $t_s^+ = 300$ (green bars) in Fig. 4.10(A), approximately 26% of the population holds anti-vaccine opinions, we find that these individuals are concentrated in even larger communities, with a higher number of these clusters (Fig. 4.10(B)). This structural shift undermines the effectiveness of all campaigns in curbing the spread of anti-vaccine influence, resulting in high epidemic sizes across all strategies (see Fig. 4.9(D)).

To gain deeper insight, we move beyond the average community size, which might mask the presence of dominant clusters and provide limited information about the underlying structural patterns. We further explore the largest existing anti-vaccine communities before the intervention. Figure 4.11 shows the average sizes of the largest and second-largest anti-vaccine communities present at the time of intervention, across varying intervention times t_s^+. We observe that up to intervention at $t_s^+ = 200$, which represents the scenario in which approximately 10% of the population has adopted the anti-vaccine opinion, the network does not exhibit any dominant cluster (see the brown, olive, and orange bars in Fig. 4.11). In these early stages of intervention, the largest clusters are comparable in size to the average cluster size, indicating no substantial deviation in the structure of existing communities (see

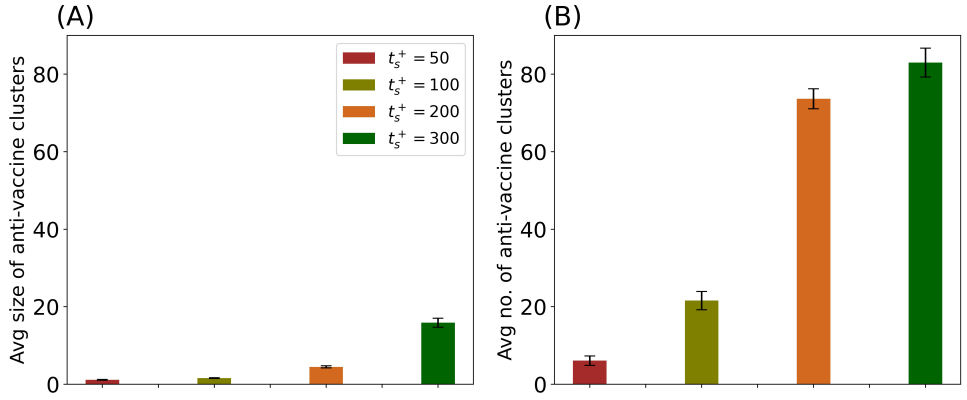

FIGURE 4.10: Analysis of anti-vaccine communities before applying the positive campaigns at different time t_s^+. Panel (A) shows the average size of existing anti-vaccine clusters, while Panel (B) displays the average number of these clusters. A total of 20 simulations were conducted. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

the brown, olive, and orange dashed lines in Figs 4.11 represent the average size of existing clusters). Therefore, intervention under such conditions can be effective, and the negative spread can be contained through the DynAdvLocT campaign.

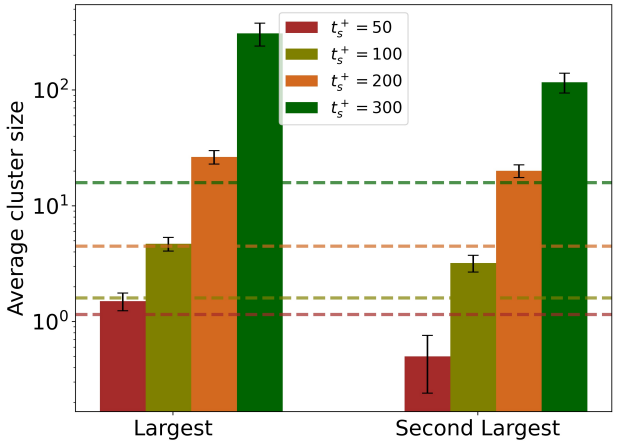

FIGURE 4.11: The average size of the largest and second-largest clusters of anti-vaccine communities before applying the intervention campaigns at different time t_s^+. Dashed lines represent the average size of existing clusters at t_s^+. A total of 20 simulations were conducted. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

However, as the negative contagion spreads further, more pronounced clustering emerges. Notably, at $t_s^+ = 300$ (green bars in Fig. 4.11), a substantial increase in cluster size is observed. The largest and second-largest anti-vaccine communities reach average sizes of 309 ± 69 and 117 ± 23, respectively, significantly exceeding the mean cluster size (green dashed line). At this point, the largest cluster comprises

approximately 6% of the population. As a result, the effectiveness of the campaigns in reducing the epidemic size is significantly limited in this scenario.

These findings support our hypothesis that the emergence of anti-vaccine communities prior to the intervention is the primary factor behind the decline in control performance across all campaigns, although the campaigns exhibit varying levels of resilience to these clusters. We find that negative propagation can be effectively contained, provided that large anti-vaccine clusters have not yet emerged. Specifically, the Min_NK campaign is more sensitive to these clusters, not because of the presence of large existing clusters, but rather due to the high number of dispersed clusters that begin to emerge from early stages, as early as $t_s^+ = 100$. Although still relatively small at this stage, they become increasingly difficult to control using the Min_NK campaign. In contrast, the DynAdvLocT campaign remains effective as long as no pre-existing giant anti-vaccine communities are present. The clustering dynamics and the growth of such communities may become even more pronounced and difficult in other network structures, which will be addressed in Chapter 6.

Beyond analyzing the anti-vaccine communities prior to the application of the campaigns, examining the impact of different campaigns on the final composition of these communities is also important. We have demonstrated the effectiveness of our campaigns with different control levels under early intervention. Therefore, to understand the impact of different campaigns on the final population composition and the formation of homogeneous communities i.e., communities that share a particular attitude toward vaccination, we visualize the network after the opinion diffusion phase showing the structure of these communities. For this purpose, we use a spatial network, specifically a 2D lattice where agents are connected to their nearest neighbors on the network [4]. We apply our campaigns and capture the network state after the opinion diffusion phase to observe the resulting cluster formations.

Figure 4.12 presents a network visualization illustrating the distribution and the structure of both anti- and pro-vaccine communities, represented by red and green colors, respectively. As shown in the figure and consistent with our earlier observations, the Min_FK campaign fails to foster positive influence in the network, resulting in the formation of a large, connected anti-vaccine cluster. The areas influenced by this campaign are limited in size and exist within predominantly negative regions, indicating the failure of this strategy to target agents capable of effectively spreading positive influence (see the green dots in Fig. 4.12(C)). In contrast, the Min_NK campaign in Fig. 4.12(D) permeates the anti-vaccine clusters and effectively curbs the formation of large clusters as it is able to seed positive influence near areas dominated by anti-vaccine opinions. A similar effect is observed with the DynAdvLocT campaign in Fig. 4.12(B). These two campaigns result in small,

[4]The use of a 2D lattice is for illustrative purposes only, as it allows for a clear visualization of the spatial distribution of homogeneous opinion communities.

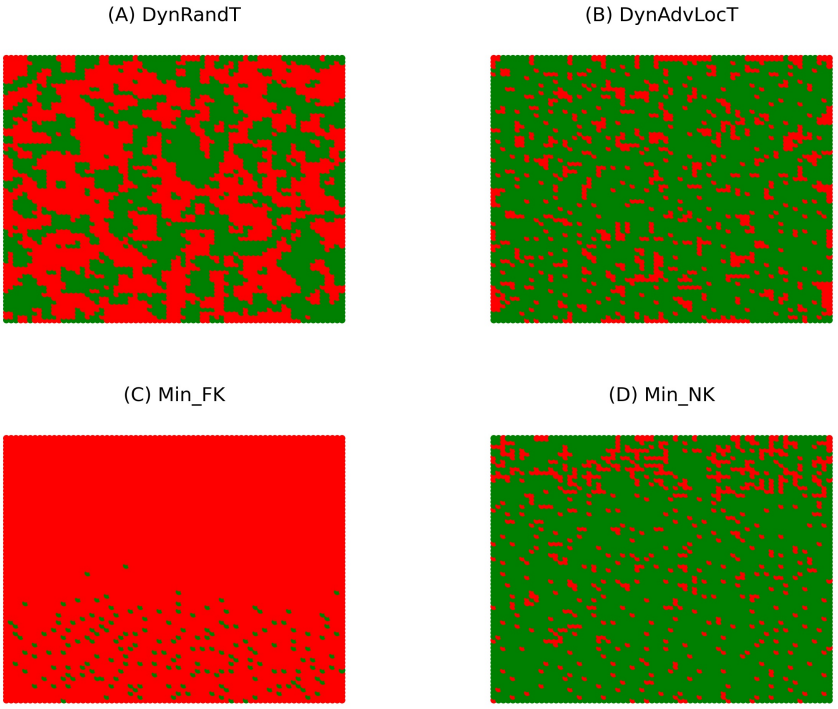

FIGURE 4.12: An example network visualization after the opinion diffusion phase, illustrating the structure of anti-vaccine communities (red) and pro-vaccine communities (green) following the implementation of different intervention campaigns. The experiments are related to a scenario where positive campaign applications occur at $t_s^+ = 100$, in 2D lattice network with $N = 5041$, $\langle k \rangle = 4$, $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

unconnected anti-vaccine clusters across the network, mitigating the growth of large-scale anti-vaccine communities. This explains the effective containment of negative diffusion achieved by these campaigns, where earlier application leads to smaller epidemic sizes. With the DynRandT campaign in Fig. 4.12(A), we observe the presence of both pro- and anti-vaccine communities, indicating its potential to foster pro-vaccine influence. However, the strategy does not effectively disrupt the connectivity of anti-vaccine clusters, as we can observe the presence of large anti-vaccine communities.

4.4 Summary

In this chapter, we advance our investigation into minimizing the spread of vaccine-preventable diseases by mitigating vaccine hesitancy through strategic

interventions under budget constraints. We formulate the problem as an optimization problem, exploring the impact of two distinct optimization objectives on reducing the final epidemic size: (1) minimizing the number of anti-vaccine opinion adopters and (2) maximizing the number of pro-vaccine opinion adopters. Accordingly, we design optimization-based campaigns that provide continuous, adaptive allocation of influence. We evaluate the effectiveness of these strategies in controlling disease spread and preventing large epidemic outbreaks, exploring which strategy contributes significantly to containing negative diffusion and strategically disrupting the connectivity among anti-vaccine agents.

Our results show that minimizing anti-vaccine opinion adoption with full knowledge of agents' social contacts and prior exposures delays the spread of negative influence in the network for an extended period; however, it ultimately fails to contain its widespread propagation, resulting in a large epidemic size. This strategy focuses on protecting individuals who are more likely to be negatively influenced and more susceptible to negative influence from their social contacts. Nevertheless, it does not effectively seed positive influence in the network, as it avoids targeting agents who have already been protected from negative influence, resulting in poor containment of negative diffusion. We then demonstrate that excluding knowledge of prior exposures enhances the campaign's effectiveness, as it enables the campaign to effectively seed and initiate positive contagion within the network, leading to more efficient control of negative diffusion and a significant reduction in the final epidemic size.

In contrast, maximizing pro-vaccine opinion adoption with full knowledge achieves better control in reducing the epidemic size than the Min_FK campaign; however, the resulting epidemic size remains significantly higher compared to the Min_NK campaign and other benchmark strategies, which demonstrate better control and a substantial reduction in epidemic size. Although this campaign increases vaccine uptake, it is an inefficient strategy to curb the spread of anti-vaccine influence. Specifically, this strategy targets agents who are most likely to be persuaded and adopt pro-vaccine opinion, leading to better seeding of positive influence and a higher number of vaccinated agents than the Min_FK campaign. However, because the campaign operates in areas distant from regions of negative influence, it promotes vaccine uptake without effectively blocking negative diffusion. As a result, anti-vaccine communities continue to expand, ultimately leading to large epidemics. Additionally, excluding historical knowledge in the maximization campaign does not improve its effectiveness; rather, it leads to poorer performance, as it causes the campaign to focus more on regions already dominated by pro-vaccine influence.

These findings suggest that minimizing the number of anti-vaccine opinion adopters leads to more effective containment of negative diffusion than strategies aimed at maximizing pro-vaccine opinion adoption. This approach is able to promote vaccination near areas of negative influence, helping to mitigate the connectivity and

growth of emerging anti-vaccine communities, resulting in a significant reduction in epidemic size. In contrast, the maximizing approach operates primarily near existing pro-vaccine regions to amplify their influence, often at a distance from negative regions, which results in ineffective control of negative diffusion and ultimately leads to a larger epidemic size. Through this analysis, we address question 1d, which explores how different optimization objectives counteract negative diffusion, influence the connectivity of existing anti-vaccine clusters, and ultimately impact the final epidemic size.

We then investigate a more realistic scenario in which interventions are initiated to combat the spread of existing anti-vaccine attitudes. Our findings reveal that our campaigns can still provide effective control even with late interventions; however, their efficacy is impacted by the extent of anti-vaccine opinion adoption in the network and the clustering structure of those adopters. We find that the minimization strategy Min_NK is highly sensitive to the presence of numerous anti-vaccine clusters, even if they are small in size. In contrast, the DynAdvLocT campaign demonstrates greater resilience to such fragmented structures and is capable of maintaining effective control, with its performance significantly reduced only in the presence of a giant anti-vaccine cluster. By doing so, we address question 2a, which investigates how intervention timing and the structural properties of existing anti-vaccine clusters affect the effectiveness of different intervention strategies.

In conclusion, early intervention can successfully contain the spread of anti-vaccine attitudes as long as large communities have not yet emerged. However, once large-scale clusters of anti-vaccine individuals have formed, containment becomes significantly more difficult. Therefore, effective management of anti-vaccine communities is crucial for minimizing the spread of an epidemic. How can we limit the growth of these communities in late intervention scenarios? In the next chapter, we extend our optimization framework and propose an optimization strategy to mitigate the expansion of anti-vaccine communities by strategically identifying critical areas in the network and targeting individuals in proximity to existing anti-vaccine clusters. To what extent can this approach contain the growth of these communities, prevent the formation of larger unprotected clusters, and reduce the epidemic size?

Chapter 5

Cluster-based Campaigns to Constrain Anti-vaccine Community Growth

Parts of this chapter have been published in the proceedings of the International Conference on Complex Networks and Their Applications (Alahmadi et al., 2025a), and have been published as a journal article in Applied Network Science (Alahmadi et al., 2025c).

5.1 Introduction

As demonstrated in the previous chapter, the timing of interventions has a significant impact on the performance of the proposed campaigns. Specifically, late interventions reduce effectiveness of intervention strategies to varying degrees. We demonstrated that this is attributed to the formation and continued growth of anti-vaccine communities, which persist despite campaign efforts to contain them. Therefore, in this chapter, we shift our focus from an individual-based intervention approach to the development of cluster-based strategies aimed at mitigating the growth and interconnectivity of emerging anti-vaccine communities.

Several studies in the existing literature emphasize the risks posed by anti-vaccine communities and their correlation with the extent of disease spread (Salathé and Bonhoeffer, 2008; Dorso et al., 2017; Gromis and Liu, 2022). Additionally, empirical research highlights the geographic clustering of unvaccinated individuals (Dubé et al., 2015), as well as the formation of distinct negative opinion clusters in online social networks (Salathé and Khandelwal, 2011). In addition, within the broader context of epidemic control, scholars have recognized the critical role of controlling network

communities in curbing disease transmission, investigating community-based immunization strategies (e.g., (Salathé and Jones, 2010; Gupta et al., 2015; Chakraborty et al., 2016; Sun et al., 2023). These studies rely on network topology to identify network communities and primarily focus on detecting and immunizing bridge agents, whose inter-community connections can facilitate disease transmission across different communities.

Here, we consider the problem of dynamic communities that emerge as a result of the propagation of anti-vaccine information, regardless of the structural properties of the underlying network. A closer examination of the literature reveals a gap in research focused on controlling the evolving dynamics of emerging anti-vaccine communities in order to mitigate their growth and connectivity. A central question we address is how to limit the expansion of these communities in the presence of ongoing negative opinion diffusion, where influence dynamics continuously reshape their structures.

To address this, we extend our optimization framework from Chapter 4 and develop cluster-based intervention strategies. These strategies aim to identify critical areas within the network and promote vaccination in those regions. Specifically, we introduce a novel method for identifying bridge (critical) nodes, defined as nodes that connect multiple anti-vaccine clusters, in dynamic settings where anti-vaccine communities evolve in response to the spread of negative influence. This approach detects these critical nodes by observing the formation of emerging anti-vaccine communities, rather than relying solely on the structural properties of the underlying network. Through our optimization framework, we aim to minimize the number of bridge nodes in the network. The goal is to strategically disrupt connections between existing communities, promote vaccine uptake in critical areas, and prevent the emergence of large-scale anti-vaccine clusters.

Next, we consider a more realistic scenario in which the available knowledge about individuals' vaccine-related opinions is subject to noise or errors. Such uncertainty is common in practice, where this information might be incomplete or imprecise. Up to this point, we have assumed access to complete and error-free knowledge. But what if the available information is incomplete or erroneous? Thus, we explore how varying levels of noise in the available information affect campaigns' performance. Understanding this impact is crucial for evaluating the robustness of our strategies and identifying which ones are more tolerant or sensitive to uncertainty and why.

In this chapter, we make the following contributions. We address research question 2b, which focuses on designing strategies to control the growth of anti-vaccine clusters and prevent the emergence of large-scale anti-vaccine communities. Specifically, we extend the optimization framework developed in Chapter 4 and propose two cluster-based intervention strategies. The first strategy aims to minimize the number of bridge agents in the network. The second strategy not only minimizes the number

of these agents but also considers the potential risk associated with each agent in the case of adopting an anti-vaccine opinion. Next, we conduct a robustness analysis to examine the impact of data noise on campaign performance. We evaluate the best-performing campaigns in the presence of noise in vaccine-related opinion data. We explore how this noise affects the effectiveness of the campaigns in curbing the spread of anti-vaccine influence and final epidemic size.

The chapter is structured as follows: Section 5.2 describes the model, formalizes the optimization problems, and introduces the intervention campaigns. Section 5.3 presents the corresponding results. Section 5.3.1 analyzes the impact of these strategies on the final epidemic size, while Section 5.3.2 explores the effect of noise on control performance and the impact on the resulting epidemic size. Finally, Section 5.4 summarizes the main findings.

5.2 Model Description and Methods

This chapter builds on the frameworks established in Chapter 3 and Chapter 4. Specifically, we consider a population of N individuals interconnected via a small-world network (Watts and Strogatz, 1998), using the same underlying models for opinion diffusion and disease spread described in Chapter 3. We extend the optimization framework from Chapter 4 and introduce cluster-based optimization campaigns aimed at mitigating the clustering of anti-vaccine individuals and preventing the formation of large anti-vaccine clusters.

5.2.1 Minimizing Bridge Agents

Here, we develop a strategy that shifts the focus from an individual-based to a cluster-based approach, aiming to identify and protect critical regions within the network. This is particularly relevant in scenarios where anti-vaccine attitudes have already become established. As demonstrated in the previous chapter, the emergence of anti-vaccine communities leads to a decline in the effectiveness of intervention strategies due to the persistent expansion of these communities and the positive budget constraints. We hypothesize that a more effective utilization of the campaign resources is to restrict the targeting space by identifying and concentrating efforts on critical areas within the network that can help limit the growth of these communities, rather than targeting the entire population. Therefore, we propose a cluster-based optimization campaign that strategically targets individuals based on their proximity to anti-vaccine communities.

To this end, we depict the network as a network of anti-vaccine clusters, with edges indicating the neutral agents that link these clusters. We define these agents as bridge

agents or critical agents. Accordingly, we construct a matrix, referred to as the cluster adjacency matrix m, where columns denote the current anti-vaccine clusters, and rows denote the candidate nodes (neutrals) within the network. A node i that is adjacent to cluster j will possess the value $m_{ij} = 1$, while $m_{ij} = 0$ otherwise. Nodes connected to more than one anti-vaccine cluster can be identified through a row-wise summation, where $\sum_j m_{ij} > 1$. The transformation of these bridge agents into a negative state will lead to the merging of existing clusters, resulting in expeditious formation of large unprotected communities within the network. A schematic representation of the network clusters and the construction of the matrix m is presented in Figure 5.1.

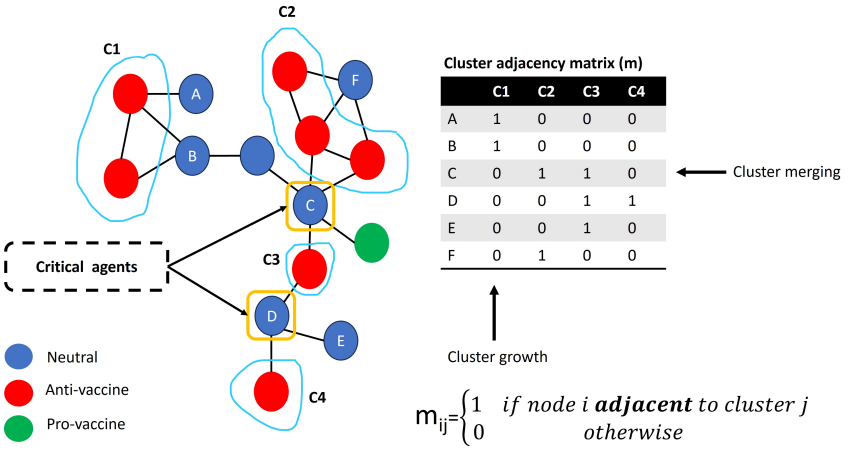

FIGURE 5.1: Visual representation of network clusters and the formation of the cluster adjacency matrix m. The left panel illustrates a network with anti-vaccine, pro-vaccine, and neutral agents, where anti-vaccine agents are distributed across multiple clusters C. Neutral agents adjacent to more than one cluster represent critical (bridge) nodes that facilitate the emergence of large communities. The right panel shows the construction of the cluster adjacency matrix m based on the network state.

Building on this, we design an objective function to minimize the number of bridge agents, rather than merely minimizing the total number of anti-vaccine opinion adopters. By targeting these nodes, we aim to enhance vaccine acceptance in critical regions near anti-vaccine communities. This campaign (denoted Cl.) is derived from the individual-based minimization campaign Min_NK outlined in the previous chapter (see Chapter 4, Section 4.2.1). By leveraging information from the cluster adjacency matrix, we can distinguish between the negative influence exerted on a given node i by neighbors belonging to the same cluster and that exerted by neighbors belonging to distinct clusters. The former is less significant when the objective is to mitigate the connectivity of anti-vaccine communities by establishing barriers between distinct clusters. Consequently, we exclude candidate targets that are adjacent to only one anti-vaccine cluster and instead focus on agents positioned between multiple anti-vaccine communities. In this regard, let N^0_{Crit} represent the set of all critical agents adjacent to more than one anti-vaccine community, i.e., $\{i \in N^0_{Crit} | \sum_j m_{ij} > 1\}$. The estimated number of critical nodes converting to a

negative state in the next time step is then given by:

$$\left\langle \Delta n^{-}_{crit} \right\rangle_{t+1} = \sum_{i \in N^{0}_{Crit}} P^{-}_{r_i}. \tag{5.1}$$

Accordingly, we can formalize our optimization problem as follows:

$$\min_{\mu^{+}} \langle \Delta n^{-}_{crit} \rangle_{t+1}, \text{ s.t.} \qquad \sum_{i} \mu^{+}_{i} \leq B, \qquad 0 \leq \mu^{+}_{i} \leq 1. \tag{5.2}$$

This represents the first cluster-based optimization campaign (Cl.).

To effectively allocate the campaign's resources, it is vital to distinguish between anti-vaccine communities based on the potential risks they pose when interconnected. Prioritizing control over larger communities is crucial to avoid exacerbating the situation by the formation of a massive, unprotected community. Therefore, we aim to further enhance this campaign by incorporating weights into the objective function, which reflect the expected size of the emergent community resulting from the conversion of a bridge node to a negative state. The weighted objective function can then be formulated as follows:

$$\left\langle \Delta n^{-}_{crit} \right\rangle_{wgt(t+1)} = \sum_{i \in N^{0}_{Crit}} (P^{-}_{r_i} . \sum_{j} (|C_j| . m_{ij})), \tag{5.3}$$

where $|C_j|$ represents the size of cluster j. Accordingly, we can formalize our optimization problem as follows:

$$\min_{\mu^{+}} \langle \Delta n^{-}_{crit} \rangle_{wgt(t+1)}, \text{ s.t.} \qquad \sum_{i} \mu^{+}_{i} \leq B, \qquad 0 \leq \mu^{+}_{i} \leq 1. \tag{5.4}$$

This represents the second cluster-based optimization campaign, referred to as (Cl.$_{wt}$).

5.3 Results

In this section, we analyze the control performance of the cluster-based optimization strategies. We begin by presenting the results for the final epidemic size in comparison to the benchmark campaigns, along with an analysis of the structure of the anti-vaccine communities (Section 5.3.1). Next, to assess the robustness of the model in the presence of uncertainty in vaccine information, we evaluate its performance across different levels of data noise (Section 5.3.2).

For all experiments, we maintain the same network settings and diffusion parameters as in Chapter 3 for comparative analysis. We follow a similar configuration to that used in Campbell and Salathé (2013), considering a small-world network (Watts and

Strogatz, 1998) with size $N = 5000$, an average degree $\langle k \rangle = 10$, and a rewiring probability $p = 0.01$. We use an opinion formation threshold of $\theta = 2$ as we consider complex contagion in which agents require more than one exposure to adopt an opinion. For the SIR model parameters, we set the infection rate $\beta = 0.1$, recovery rate $\gamma = 0.1$, and seed set $I_0 = 1$. In all figures presented, the error bars correspond to the 95% confidence intervals. Unless otherwise stated, the results for each scenario are averaged over 50 realizations, each of which involves generating a network and simulating the opinion diffusion process followed by the disease spread. Furthermore, for the optimization campaigns if the budget B is not fully utilized by the campaign, we randomly choose T nodes and distribute the remaining budget equally among them.

5.3.1 Results of Minimizing Bridge Agents

To evaluate the control performance and assess the effectiveness of the cluster-based campaigns in mitigating negative influence, we present the resulting epidemic size when the intervention is applied at different times t_s^+. Additionally, we present the final number of anti-vaccine individuals and the average number of resulting anti-vaccine communities. We benchmark its effectiveness against the best-performing campaigns proposed in previous chapters, specifically DynAdvLocT and Min_NK, as well as the random campaign, DynRandT as a baseline strategy. Fig. 5.2 presents the results compared to the benchmark campaigns, where Figs.5.2 (A)-(C) display the final epidemic size, (D)-(F) show the final number of anti-vaccine agents, and (G)-(I) illustrate the final number of anti-vaccine clusters.

As shown in Figs. 5.2 (A)-(C), the cluster-based approach (Cl. and Cl.$_{wt}$ campaigns, represented by the red and olive bars) demonstrates a robust capacity to contain the propagation of negative influence during late interventions, resulting in a smaller epidemic size compared to other strategies. Notably, these campaigns remain effective even when applied as late as $t_s^+ = 300$ (Fig. 5.2 (C))—a stage where approximately 26% of the population is already negatively influenced and large anti-vaccine communities have formed (see green bars of Fig. 4.11). In addition, the weighted campaign Cl.$_{wt}$ (olive bars), which accounts for the potential risk associated with each bridge node, demonstrates greater effectiveness than the unweighted cluster-based campaign Cl. (red bars). These findings indicate that strategically targeting bridge nodes within the network contributes effectively in promoting vaccine uptake in critical areas and establishing barriers between existing anti-vaccine communities. This limits cluster expansion and prevents the formation of large unprotected communities, significantly curtails disease transmission, and results in smaller epidemic sizes than benchmark campaigns.

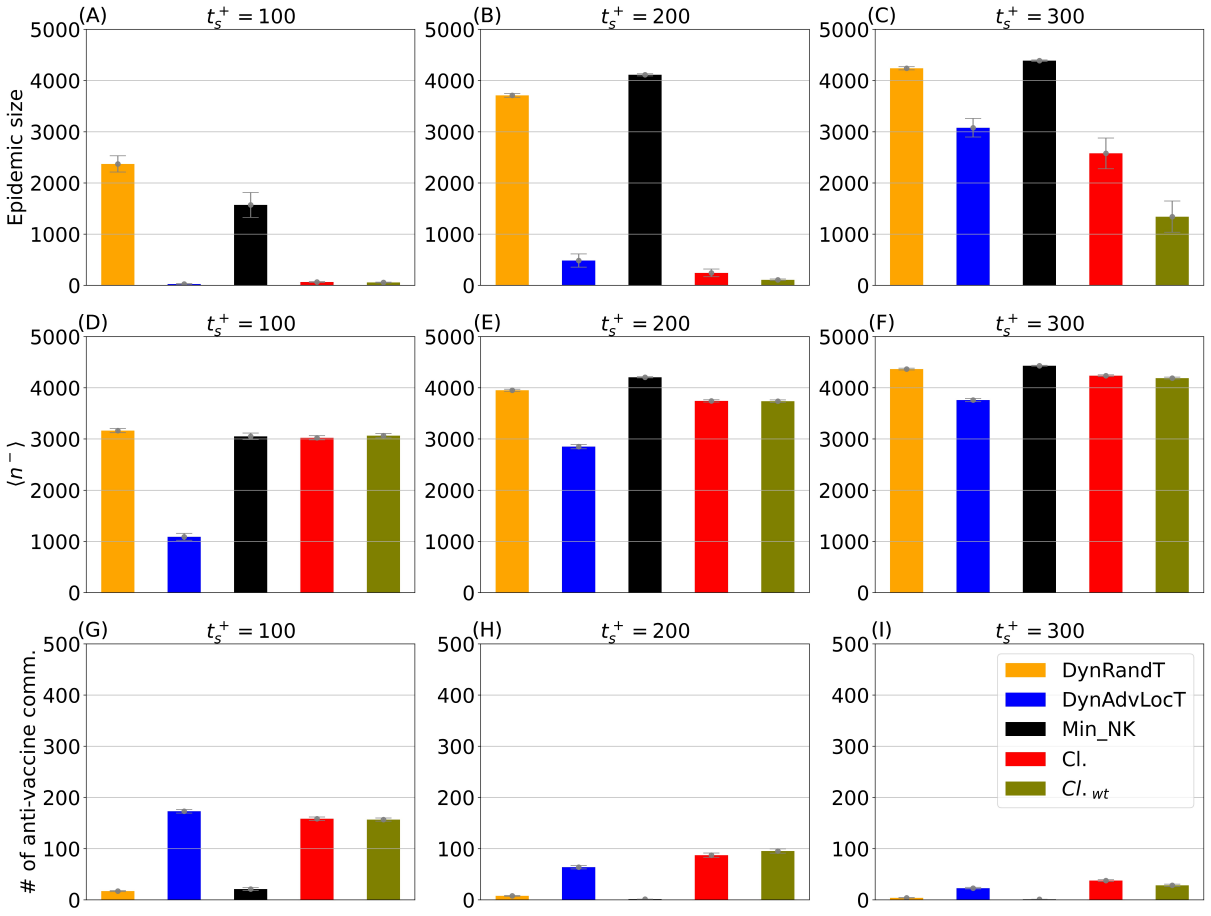

FIGURE 5.2: Results of cluster-based campaigns (Cl. and Cl.$_{wt}$) compared to benchmark campaigns: DynAdvLocT, Min_NK, and DynRandT, with campaigns applied at intervention time t_s^+. (A)-(C) Final epidemic size, (D)-(F) final number of anti-vaccine agents, and (G)-(I) final number of anti-vaccine communities. The experimental settings are: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and T=5.

To further understand the impact of this approach on mitigating the connectivity of anti-vaccine communities, we present the number of anti-vaccine opinion adopters and the average number of final anti-vaccine clusters in which these agents are distributed in Figs. 5.2(D)–(I). As shown in the figure, the cluster-based approach does not substantially minimize the total number of anti-vaccine individuals compared to the Min_NK campaign (Fig. 5.2(D)-(F)). However, it results in a higher number of final anti-vaccine communities, indicating that approximately the same number of anti-vaccine agents are distributed across more clusters with the cluster-based campaigns (red and olive bars in Figs. 5.2(G)-(I)). In contrast, these individuals tend to concentrate within fewer communities with the Min_NK campaign (black bars in

Figs.5.2(G)-(I)). This demonstrates that the cluster-based approach is effective in mitigating both the connectivity between existing anti-vaccine communities and the formation of large clusters. As a result, these strategies lead to a notable reduction in epidemic size, even when interventions are deployed at a late stage (e.g., after approximately 26% of the population has already adopted anti-vaccine opinions).

5.3.2 Addressing Vaccine-Related Information Uncertainty

In this section, we address the uncertainty associated with the available data regarding individuals' vaccination attitudes. As our intervention strategies primarily rely on information about individuals' vaccine-related opinions, one limitation of our model lies in the resolution of the available data on these attitudes. Up to this point, we have assumed complete knowledge of these opinions to gain a deeper understanding of how different interventions perform. However, in practice, this assumption may not accurately reflect real-world scenarios, where such information is often incomplete or subject to error. We refer to these inaccuracies as noise or error in the data. Therefore, it is crucial to account for this uncertainty to evaluate the robustness and reliability of our model.

To this end, we introduce an error parameter, ϵ, which represents the rate of error (or noise) in the available data, and conduct a parameter sweep to explore its impact. We aim to examine how different levels of error affect control performance, allowing us to better understand the impact of these uncertainties across different interventions and identify which strategies are more tolerant to such errors. To model this uncertainty, each agent in the network is randomly assigned one of the three possible states, neutral, anti-vaccine, or pro-vaccine, with probability ϵ. Specifically, at each time step t, the model operates as follows for each agent i: with probability $1-\epsilon$, the agent's state s_i is assumed to be correct and remains unchanged; with probability ϵ, the state is assumed to be unknown, and the agent is assigned a randomly selected state from the three possible states in the system. Subsequently, campaign resources are allocated based on the noisy network, referred to as $\hat{G}$, while the diffusion of vaccine opinions occurs within the actual network G.

5.3.2.1 Impact of Noise on Control Performance

To evaluate the impact of data noise, we present the dependency of the final epidemic size on different levels of noise, ϵ, in Fig. 5.3. We consider the baseline DynRandT campaign (Fig. 5.3(A)), the best-performing campaigns DynAdvLocT (Fig. 5.3(B)) and Cl._{wt} (Fig. 5.3(D)), as well as the Min_NK campaign (Fig. 5.3(C)), due to its effective control with early intervention. We also consider three distinct budget scenarios: (i) a budget disadvantage scenario (dotted lines), where the opponent has a larger budget

($\mu^+ < \mu^-$); (ii) an equal-budget scenario (dashed lines), where the opponent and the positive campaign have equal budgets ($\mu^+ = \mu^-$); and (iii) a budget advantage scenario (solid lines), where the opponent has a smaller budget than the positive campaign ($\mu^+ > \mu^-$). In addition, the findings pertain to the intervention time scenario at $t_s^+ = 100$, an intervention time where campaigns can still provide a notable control.

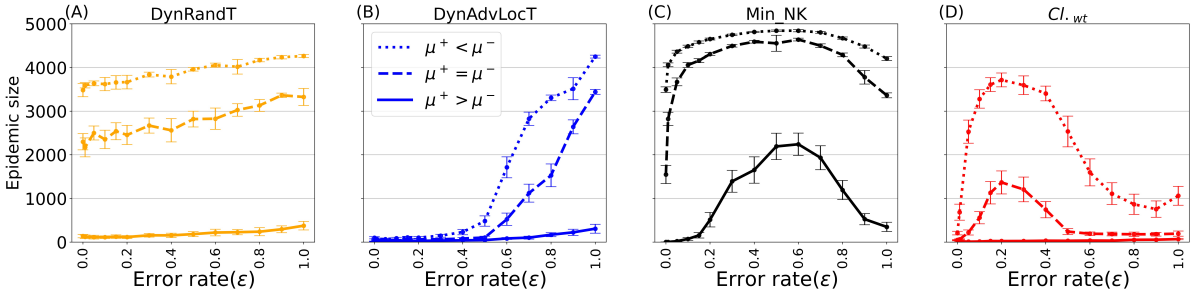

FIGURE 5.3: Final epidemic size for different noise levels ϵ, considering the best intervention campaigns. Applying campaigns at $t_s^+ = 100$ with approximately 1% of the population had adopted the anti-vaccine opinion. The experimental settings are: $\mu^- = 0.001$, $\omega = 0.006$, and $T = 5$ except for $\mu^+ > \mu^-$ scenario $T = 10$ for the DynRandT and DynAdvLocT heuristics.

From Fig. 5.3, we can make several key observations. First, a low noise rate (i.e., $\epsilon \leq 0.1$) does not significantly affect the level of control achieved by most campaigns compared to their performance in the noise-free scenario ($\epsilon = 0.0$). An exception occurs in the disadvantage scenarios for the Min_NK and Cl._wt campaigns and equal budget scenario for the Min_NK campaign, where performance degrades even with low noise conditions. This threshold extends as the positive budget increases. Second, heuristic campaigns (DynRandT and DynAdvLocT) show a steady increase in epidemic size as error increases. Optimization-based campaigns initially exhibit a similar pattern with significant performance degradation. However, their sensitivity is reduced and control performance progressively improves with higher noise levels, resulting in better outcomes at higher error rates compared to lower ones. Third, the campaigns exhibit varying levels of sensitivity to noise. Below, we describe in detail the behavior of each campaign in response to noise.

For the DynRandT campaign (Fig. 5.3(A)), the results demonstrate relatively stable performance, with only a slight and gradual decline as the error rate increases. Moreover, the effectiveness of this campaign improves with a higher budget, as shown in the budget advantage scenario represented by the solid line in Fig. 5.3(A), and continues to provide similarly effective control even under high levels of noise, with only a slight increase in epidemic size at very high noise levels (see $\epsilon = 1.0$). Thus, the campaign is not significantly affected by data noise. However, its overall performance is relatively poor compared to other campaigns, particularly in budget disadvantage and equal-budget scenarios.

For the DynAdvLocT campaign (Fig. 5.3(B)), the results demonstrate a high tolerance to noise up to $\epsilon = 0.5$, even under the budget disadvantage scenario (dotted line). Beyond this point, performance gradually declines and reaches its lowest levels under very high noise conditions ($\epsilon \geq 0.8$) in both the budget disadvantage and equal-budget scenarios (dashed and dotted lines of Fig. 5.3(B)). In contrast, under the budget advantage scenario, the campaign maintains effective control and exhibits low sensitivity to noise (solid line of Fig. 5.3(B)).

For the Min_NK campaign (Fig. 5.3(C)), the results indicate high sensitivity to noise, with performance degrading significantly even at a low noise rate (i.e., $\epsilon = 0.1$). This effect is particularly pronounced in the budget disadvantage and equal-budget scenarios (dotted and dashed lines in Fig. 5.3(C)). With an increased budget in the budget advantage scenario (solid line), we observe that the campaign's tolerance improves, with the decline in performance beginning at a higher noise level (i.e., $\epsilon > 0.2$). Additionally, when noise exceeds 50% (i.e., $\epsilon > 0.5$), the tolerance of the campaign increases and performance gradually improves, leading to a steady decrease in epidemic size across all budget scenarios.

The improved tolerance of the Min_NK campaign in high-noise scenarios can be attributed to changes in agents' neighborhood structures resulting from noise. In such scenarios, the noise benefits the campaign by pushing certain nodes into more favorable states that meet the campaign's targeting criteria, whereas they were less likely to be selected. These nodes become more visible and suitable candidates for the campaign due to alterations in their neighborhood structures caused by noise. The emergence of these new nodes allows the campaign to expand its reach in promoting the pro-vaccine influence in the network, which compensates for distortions that occur at lower noise levels and hence enhances its performance. We will discuss this in detail in the following section.

For the cluster-based Cl.$_{wt}$ campaign (Fig. 5.3(D)), the results demonstrate resilience in maintaining robust control even at high noise levels across most scenarios, with notable degradation occurring only at lower noise levels. Furthermore, it shows high sensitivity to noise only in the budget disadvantage scenario, although its performance improves as the noise level increases. Specifically, we find that performance declines at low to moderate noise levels in both the equal-budget and budget disadvantage scenarios, approximately when $0.1 < \epsilon < 0.4$, with the range being slightly larger in the budget disadvantage case. Despite this degradation, we observe that beyond a certain threshold (approximately $\epsilon \geq 0.2$), performance begins to improve as the noise level increases, eventually stabilizing and achieving effective control for $\epsilon > 0.5$. This behavior is similar to that observed in the Min_NK campaign, although the Cl.$_{wt}$ campaign achieves substantially better control performance. In the budget advantage scenario, the campaign consistently provides robust performance and exhibits low sensitivity to noise.

In the following section, we discuss in detail why optimization campaigns become less sensitive to noise as the noise level increases. We also show how noise impacts the social network structure of agents, specifically, their neighborhood configurations.

5.3.2.2 Impact of Noise on Social Network Structure

Our findings show that optimization campaigns (Min_NK and Cl_{wt}) exhibit fluctuations in their sensitivity to noise in which they are more sensitive at low noise levels, but become less sensitive as noise increases (Figs.5.3(C) and (D)). As one might expect, the presence of noise alters the social network structure (the neighborhood configurations) of agents within the system. This, in turn, affects the performance of optimization campaigns, as these campaigns target agents based on their social network states, prioritizing those surrounded by many anti-vaccine neighbors. However, why do optimization campaigns exhibit varying levels of sensitivity under different noise levels? In the following, we provide an analysis of how varying noise rates impact the social network structure of agents within the network.

The first question that arises is: why does a higher noise rate reduce the sensitivity of optimization campaigns to noise and enhance their resilience, leading to a smaller epidemic size compared to lower noise levels? As mentioned earlier, the improved tolerance of optimization campaigns can be attributed to changes in agents' neighborhood structures, which bring more nodes into focus by meeting the campaign's selection criteria and hence increase their likelihood of being targeted.

To demonstrate this further, we present the real neighborhood structure (noise-free network) of neutral nodes, including those that were targeted based on the noisy neighborhood structure (noisy network). Figure 5.4 depicts the neighborhood structures of neutral agents in the absence of noise in Figs. 5.4(A)-(E) (upper panels) and in the presence of noise in Figs. 5.4(F)-(J) (lower panels) for the Min_NK campaign. We aim to compare the social network structures of targeted individuals in noisy and noise-free networks. For the noise-free neighborhood structures, we use the neighboring information of the neutral agents from the original network G, while for the noisy neighborhood structures, we use the corresponding information from the noisy network $\hat{G}$.

As shown earlier in Fig. 5.3(C), the campaign's performance begins to improve at noise levels above $\epsilon > 0.5$. By examining the neighborhood structures of the targeted agents in Fig.5.4, specifically in the panels for $\epsilon > 0.5$, we find that, in the absence of noise (Figs. 5.4(D) and (E)), these agents have relatively few anti-vaccine neighbors and were thus less likely to be selected by the Min_NK campaign. However, with high noise rates (Figs. 5.4(I) and (J)), these agents acquire a greater number of apparently anti-vaccine neighbors, making them meet the targeting criteria and increasing their

likelihood of being targeted by the campaigns (see their number of anti-vaccine neighbors in Figs.5.4 (I) and (J)).

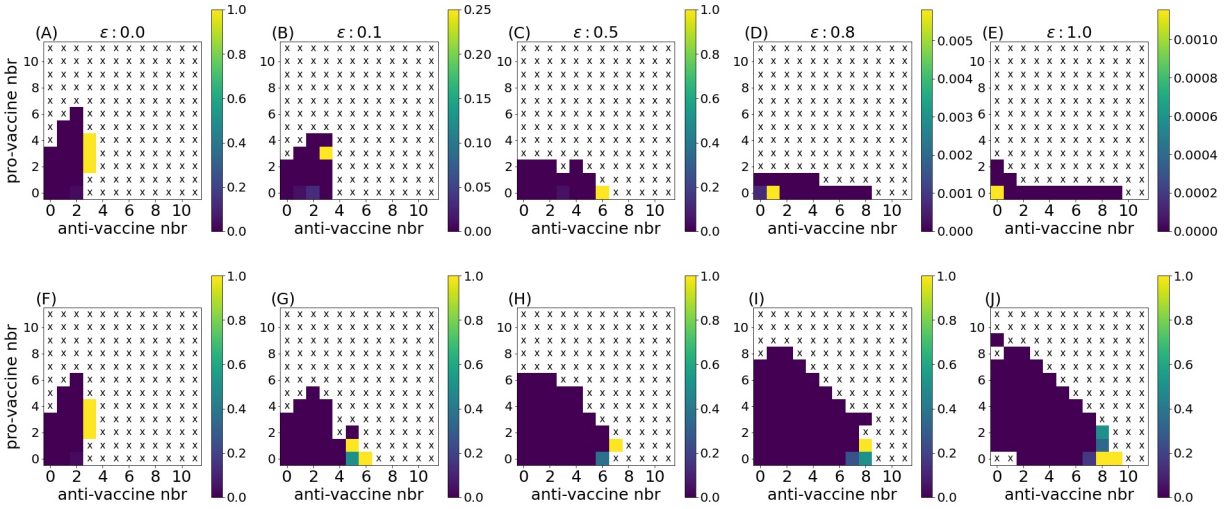

FIGURE 5.4: An example of neighborhood structures for uncommitted neutral agents, displaying the number of pro-vaccine and anti-vaccine neighbors. The figure shows the structures without noise (upper panels) and with noise (lower panels) at different error rates ϵ at time $t = 150$. Each cell represents the average allocation for agents with x anti-vaccine neighbors and y pro-vaccine neighbors. An 'x' in the figure further denotes that no agent was found in this particular neighborhood configuration. Targeted agents are those with a non-zero allocation. Applying the Min_NK campaign at $t_s^+ = 100$ with an average of 35 existing anti-vaccine adopters, approximately 0.01% of the population. The experiment setting are: $\mu^+ = 0.001, \mu^- = 0.001$, and $\omega = 0.006$, $T = 5$.

In addition, these targets have few or no pro-vaccine neighbors in the noise-free network (Figs. 5.4(D) and (E)), implying that most of their social connections are neutral agents. As previously demonstrated in Chapter 3 (see specifically local information based campaigns in Section 3.3.2.2), targeting agents with anti-vaccine neighbors while also having neutral neighbors is an effective approach for curbing negative diffusion and facilitating the spread of positive influence. These agents are more effective in promoting positive influence through their neutral networks which amplify the campaign's reach. This enables the optimization campaigns to regain control and exhibit low sensitivity to noise as the noise level increases. The same reasoning applies to the cluster-based campaign (Cl.$_{wt}$). The emergence of these new nodes with high noise levels also broadens the targeting pool for this campaign which compensates for the distortions that occur at lower noise levels, thereby reducing the campaign's sensitivity to noise and allowing it to maintain its effectiveness.

To generalize our findings, we compare the average number of anti-vaccine neighbors for the targeted agents in both noisy and noise-free networks. Specifically, we

compute the average number of anti-vaccine neighbors in the noisy network, denoted by $\hat{K_i^-}$, and in the noise-free network, K_i^- to quantify the difference for targeted agents. Fig.5.5 presents the average number of anti-vaccine neighbors among the target sets selected by the Min_NK campaign (Fig. 5.5(A)) and the Cl.$_{wt}$ campaign (Fig. 5.5(B)). Results corresponding to the noise-free network are shown in red, while those from the noisy network are shown in green. In this analysis, we capture both the noise-free network G and the noisy $\hat{G}$ at time step t for different error levels ϵ, and compute the averages at that time step.

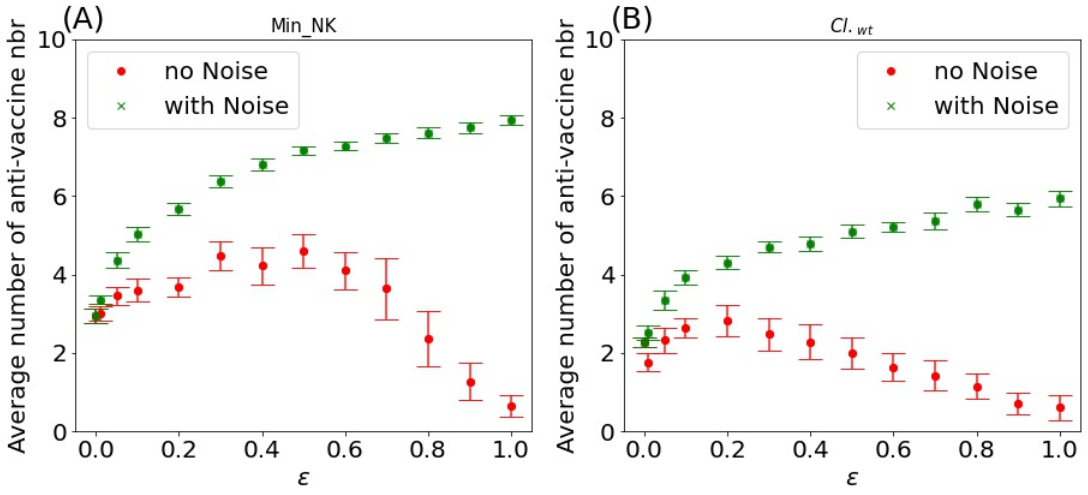

FIGURE 5.5: Average number of anti-vaccine neighbors for the targeted agents by (A) the Min_NK campain and (B) the Cl.$_{wt}$ campaign, for different error rates ϵ. Campaigns are applied at $t_s^+ = 100$ with approximately 1% of the population had adopted the anti-vaccine opinion. Results are based on networks captured at $t = 200$ and averaged over 30 simulations, with parameters $\mu^+ = 0.001, \mu^- = 0.001, \omega = 0.006$.

From Fig. 5.5, we observe that the agents targeted by the Min_NK campaign (Fig. 5.5(A)) and the Cl.$_{wt}$ campaign (Fig. 5.5(B)) in the noisy network (green points) are selected as agents having a high number of anti-vaccine neighbors. However, analyzing their neighborhood structures based on the original noise-free network G (red points) reveals that these agents in fact have fewer anti-vaccine neighbors than they appear to under noisy conditions. This difference becomes more pronounced as ϵ increases, clearly indicating that higher noise levels lead to structural alterations that bring more neutral agents into focus for the campaign. This supports the observation that more nodes become noticeable to the campaigns, thereby increasing their likelihood of being targeted due to changes in their neighborhood structure.

A second question that emerges is why this effect is not observed at lower noise levels. At low to moderate noise levels, noise results in performance degradation. This is because the noise at these levels introduces subtle yet significant alterations to agents' social networks that primarily affect structurally similar nodes. In other words, the targeted agents exhibit similar neighborhood structures in both the noisy and noise-free networks; however, the presence of noise introduces distortions that affect

the optimal selection among them. As a result, the campaign may target agents with similar neighborhood structures, but these agents are not the optimal choices, such as neutral targets located near pro-vaccine neighbors. For instance, in Figs. 5.4(B) and (G), specifically at $\epsilon = 0.1$, there are only minor differences in the neighborhood structures of the targeted agents in the noisy (Fig.5.4(B)) and noise-free networks (Fig.5.4(G)). The number of anti-vaccine neighbors among the targeted agents (non-purple cells) is nearly identical, around five in the noisy network compared to three in reality. However, in reality, these agents are located near pro-vaccine neighbors, which diminishes the effectiveness of the targeting scheme, as they are not optimal choices. This is also observed in Fig.5.5, where the number of anti-vaccine neighbors in the noisy and noise-free networks is relatively close for low to moderate noise levels (i.e., $\epsilon \leq 0.5$ in Fig.5.5(A) and $\epsilon \leq 0.4$ in Fig.5.5(B)).

5.4 Summary

In this chapter, we propose an optimization-based intervention strategy aimed at mitigating the growth and connectivity of existing anti-vaccine communities. This campaign specifically targets and minimizes bridge agents, those who connect distinct communities, thereby disrupting connectivity that contributes to the expansion of anti-vaccine clusters. As a result, the strategy effectively identifies and protects critical network areas, limiting the formation of large anti-vaccine communities and leading to a significant reduction in the resulting epidemic size. We further optimize this strategy by incorporating knowledge about the potential risk associated with each bridge node. This knowledge reflects the expected size of the emergent anti-vaccine community that may arise as a result of bridge node conversion. This approach is even more effective in controlling negative diffusion and leads to a more substantial reduction in epidemic size with the same budget constraints. Through this work, we address research question 2b, which investigates how to design a strategy to mitigate the interconnectivity and expansion of anti-vaccine clusters, and to what extent this reduces the resulting epidemic size.

Next, since our intervention strategies heavily depend on information about individuals' opinions, investigating the impact of potential noise in the information on control performance is crucial. To this end, we evaluated the impact of noise on the effectiveness of the proposed strategies. Our results reveal that campaigns' tolerance to noise increases as the available budget increases. In the budget advantage scenario, almost all campaigns (except for the Min_NK campaign) maintained efficient performance, exhibiting only slight performance degradation at high levels of noise. The impact of noise was more pronounced in the equal-budget and budget disadvantage scenarios. As one might expect, in these cases, control performance

declined as noise increased, however we observed varying levels of sensitivity across campaigns.

We find that the performance of the DynRandT and DynAdvLocT campaigns decreases as noise levels increase. In contrast, optimization campaigns (Min_NK and Cl_wt) exhibit a different pattern, being more sensitive to low levels of noise and less affected by high levels of noise. Specifically, the DynRandT campaign is less sensitive to noise, with only a slight decrease in performance as noise levels increase. The DynAdvLocT campaign is the most resilient across different budget scenarios, with performance degradation beginning only at noise levels above 50%. On the other hand, we find that the performance of the Min_NK campaign degrades significantly with the presence of noise. In contrast, the Cl_wt campaign demonstrates greater robustness, with its worst performance observed only under the budget disadvantage scenario at low to moderate noise levels.

We demonstrate that, for optimization campaigns, the system becomes too chaotic at low to intermediate noise levels for the campaigns to maintain effective control, resulting in performance degradation as the noise disrupts their ability to select the optimal candidates among structurally similar agents. Interestingly, at high noise levels, performance improves from its lowest point, and the campaigns' tolerance to noise increases. This occurs because high noise levels bring more agents into focus, increasing their likelihood of being targeted. These nodes, due to their neutral social connections, help amplify positive contagion within the network, thereby enhancing the campaign's resilience to noise. These results suggest that using a deliberately noisy strategy might be beneficial. By doing so, we address a key practical challenge related to the potential presence of noise in real-world data and its impact on the performance of intervention strategies. Specifically, which strategies are more resilient, which are more sensitive, and the underlying reasons for these differences.

In conclusion, we proposed a novel strategy that identifies and protects critical areas within the network. As demonstrated, this strategy provides efficient control and exhibits robustness under late intervention, as well as a high tolerance to noise. A natural question that arises is whether this strategy remains optimal when applied to different network structures. Thus far, we have used a small-world network, which exhibits a homogeneous structure, but what if we consider heterogeneous structures? Are the patterns of growth and formation of anti-vaccine communities in such networks comparable to those observed in the topology considered so far? Additionally, we demonstrated the effectiveness and robustness of the DynAdvLocT campaign. Will this strategy remain efficient across different network topologies?

In the following chapter, we investigate how variations in network topology influence the performance of intervention campaigns and explore potential strategies to further improve our campaigns.

Chapter 6

Impact of Network Structure

Parts of this chapter have been published as a journal article in Applied Network Science (Alahmadi et al., 2025c).

6.1 Introduction

So far in the thesis, social interactions have been modeled as a small-world networks. While this structure has been exhibited in many real-world social networks, other structural patterns have also been observed. Therefore, accounting for diverse network topologies and their structural dynamics is crucial, as network structure plays a significant role in shaping information diffusion. A robust understanding of how anti-vaccine opinions propagate and anti-vaccine communities emerge in different network topologies is essential for assessing their impact on public health campaigns.

Networks, whether representing social interactions or disease transmission pathways, exhibit considerable structural variability arising from differences in centrality, clustering, connectivity patterns, and other topological metrics (Boccaletti et al., 2006; Newman, 2018). This influences how information, behaviors, or pathogens spread (Hein et al., 2006; Varga, 2017; Pérez-Ortiz et al., 2022). Although dynamic processes on networks have attracted considerable attention across multiple areas such as the influence of network structure on diffusion processes (Pérez-Ortiz et al., 2022), influence maximization methods (Erkol et al., 2019), and interventions on networks (e.g., (Shi et al., 2020); for a review, see (Riddell et al., 2025)), there is a lack of systematic understanding of how different network topologies affect mitigation efforts and influence the effectiveness of various intervention strategies, particularly in coupled diffusion processes. A campaign strategy that works well in one network type may fail in another. This variability underscores the need to study how different intervention strategies perform across network structures. By exploring diverse

campaigns in different network structures, we aim to uncover the interplay between network topology and intervention effectiveness, and how each control strategy is influenced by the underlying structure.

In this chapter, we extend our work by applying the proposed intervention strategies to more general classes of networks to assess whether the gains in control previously observed in small-world networks persist across various topologies. In addition to the small-world networks discussed earlier in this thesis, we also examine five more topologies, including scale-free networks, which have also been observed in social networks (Barabási and Albert, 1999; Mislove et al., 2007; Sadikov and Martinez, 2009), as well as random networks, regular lattices, and real-world social networks. Additionally, given the varying heterogeneity in degree distribution across the diverse network structures considered, we extend our optimization framework by developing a new strategy that explicitly account for this heterogeneity.

Specifically, in this chapter, we address research questions 3a and 3b detailed in Section 1.1 through the following contributions. We systematically explore the impact of network topologies on the control performance of the proposed campaigns and identify the most effective strategy for each structure, and analyze the underlying reasons for these outcomes. We also study how the formation of anti-vaccine communities varies across these topologies at different intervention timings and how their structures impact containment efforts. Then, we evaluate the sensitivity of control performance to the campaign's budget across these networks.

The chapter is structured as follows. Section 6.2 describes the model, including the new strategy and the network structures considered. Then, Section 6.3 presents the results and analysis. Specifically, Section 6.3.1 presents the results for the epidemic size obtained across different networks. Section 6.3.2 analyzes the clustering patterns and formation of anti-vaccine communities. Section 6.3.3 examines the influence of budget size. Section 6.3.4 assesses performance on real-world networks. Finally, Section 6.4 summarizes the main findings.

6.2 Model Description and Methods

As before, we consider a population of N individuals interacting within a social network. This time, however, they are connected through different network structures. In addition to the small-world network (SW) discussed earlier in this thesis, we also consider Erdős–Rényi random graphs (RN) (Erdős and Rényi, 1959), Barabasi-Albert scale-free networks (SF) (Barabási and Albert, 1999), regular lattices (referred to as spatial networks (SpN)), and two real-world networks—socfb-Maine59 and socfb-JohnsHopkins55—extracted from the Facebook platform (Rossi and Ahmed, 2015). The basic properties of all the networks considered are summarized in

Network type	Network	Size	Avg deg	Max deg	Clust coeff	Density
Real-world	JohnsHopkins55	5.2K	72	886	0.267	0.014
	Maine59	9.1K	53	1K	0.239	0.006
Synthetic	Small-world (SW) (Watts-Strogatz)	5K	10	12	0.65	0.002
	Scale-free (SF) (Barabási–Albert)	5K	10	367	0.013	0.002
	Random graph (RN) (Erdös-Rényi)	5K	10	24	0.002	0.002
	Regular lattice (SpN)	5K	10	10	0.67	0.002

TABLE 6.1: Basic properties of the networks considered, including network size, average degree (Avg deg), maximum degree (Max deg), clustering coefficient (Clust coeff), and density, which were previously explained in Chapter 2

Table 6.1. Opinion formation and propagation as well as the disease transmission follow the model proposed in Chapter 3.

Real social networks exhibit complex and diverse structural properties, such as small-world properties and degree heterogeneity, that cannot be fully captured by a single network model. Here, we consider a range of network types that span different structural characteristics. These networks vary in clustering, degree distributions, and connectivity patterns, ranging from highly clustered small-world and spatial networks to highly degree-heterogeneous scale-free and random Erdős–Rényi networks, while real-world networks provide insights into diffusion dynamics and control in realistic network structures.

Scale-free networks and small-world networks have been observed in many real-world networks (Watts and Strogatz, 1998; Barabási and Albert, 1999; Barrat and Weigt, 2000; Mislove et al., 2007; Sadikov and Martinez, 2009). Random networks, although they lack inherent structure and do not accurately reflect the structure of real-world networks, are well-known and widely used in the literature (e.g., (Salathé and Bonhoeffer, 2008; Pérez-Ortiz et al., 2022)). They provide essential baseline models to strengthen our understanding of how campaigns perform in unstructured networks. Regular lattices, on the other hand, capture local interactions constrained by spatial proximity, reflecting communication with immediate neighbors with the absence of long-range connections.

A significant factor introduced in this chapter, particularly through the use of scale-free networks, is heterogeneity of degree distribution. This heterogeneity serves as a distinguishing characteristic between different network types. Scale-free networks are defined by the presence of hub nodes—nodes with a high number of connections and represents a small fraction of the overall population. These hubs play a crucial role in accelerating the diffusion of spreadable entities (Dezső and Barabási, 2002; Hein et al., 2006). The study of diffusion processes and control in heterogeneous topologies

is well-established in the literature (Dezső and Barabási, 2002; Sood and Redner, 2005; Zhang et al., 2010; Budak et al., 2011; Masuda, 2015; Montes et al., 2020). Targeting hub nodes in such networks is a commonly used control heuristic in the literature, applied to different diffusion processes such as epidemic control (Dezső and Barabási, 2002), immunization campaigns (Zhang et al., 2017; Prieto Curiel and González Ramírez, 2021), and the influence maximization problem (Kempe et al., 2003; Chen et al., 2009; Masuda, 2015; Erkol et al., 2019). Building on this, in this section, we aim to further improve our control framework by developing a new campaign that explicitly incorporates this heterogeneity and accounts for the roles of high-degree nodes.

To this end, we develop a new strategy that builds upon the Min_NK campaign proposed in Chapter 4, which aims to minimize the number of anti-vaccine opinion adopters. We extend this campaign by introducing a weight n_i^0, representing the number of neutral neighbors an agent i has. Building on the objective function in Eq. 4.12, defined in Section 4.2.1 to estimate the expected number of anti-vaccine opinion adopters, we define a weighted objective function that accounts for the particular role of high-degree nodes as follows:

$$\langle \Delta n^- \rangle_{wgt(t+1)} = \sum_{i \in N^0} Pr_i^- n_i^0. \tag{6.1}$$

Our goal is to minimize this function, incorporating weights assigned to individual agents. The budget-constraint optimization problem can therefore be formalized with respect to $\boldsymbol{\mu}^+$ as follows:

$$\min_{\boldsymbol{\mu}^+} \langle \Delta n^- \rangle_{wgt(t+1)}, \text{ s.t.} \qquad \sum_i \mu_i^+ \leq B, \qquad 0 \leq \mu_i^+ \leq 1. \tag{6.2}$$

We refer to this as the weight-augmented individual campaign (Min_NK$_{wt}$). By weighting nodes based on the number of neutral neighbors, the campaign implicitly accounts for the role of high-degree nodes in heterogeneous networks. This is because nodes with many neutral neighbors often correspond to high-degree (hubs) due to their high connectivity, particularly in the early stages of intervention when these hubs are more likely to remain unaffected by negative influence. In addition, at this stage, as negative influence has not yet widely propagated, the number of neutral neighbors is expected to closely align with a node's degree. Incorporating knowledge about neutral neighbors, rather than solely relying on node degree, provides a strategic opportunity for positive influence to spread more widely across the network. Since these neutral neighbors have not yet formed a vaccine-related opinion, they are more susceptible to positive influence.

6.3 Results

We first evaluate the control effectiveness of the intervention campaigns across different network topologies (Section 6.3.1). Next, we explore how the formation of anti-vaccine communities varies across these topologies (Section 6.3.2). Subsequently, we analyze the sensitivity of control performance to the positive budget size (Section 6.3.3). Then, we evaluate the performance in real-world networks to assess whether the control achieved is consistent with the insights gained from the dynamics of synthetic networks (Section 6.3.4).

In this chapter, we also evaluate the campaigns at different intervention timings across various network structures. Unless otherwise stated, we consider intervention timings of $t^+ = 100$, $t^+ = 175$, and $t^+ = 200$. We start from $t^+ = 100$ and progressively extend to more challenging scenarios in which negative influence spreads more widely, continuing until campaigns performance decline significantly, i.e., at $t^+ = 200$, where the majority of the population has been infected, particularly in the newly introduced network structures including scale-free and random networks. Considering scenarios beyond this point, such as $t^+ = 300$, provide no further insights into campaigns performance and are therefore excluded.

For all experiments, we applied the same network structure and diffusion parameters settings used in previous chapters. Specifically, regarding the network settings for synthetic networks, we follow a similar configuration to that used in Campbell and Salathé (2013), considering a network with size $N = 5000$, an average degree $\langle k \rangle = 10$, and a rewiring probability $p = 0.01$ for the small-world-network. We maintained a constant average degree across all synthetic networks for comparison purposes. For all networks, we used the same setting of diffusion parameters. The opinion formation threshold is set to $\theta = 2$. The SIR model parameters are set as follows: an infection rate $\beta = 0.1$, a recovery rate $\gamma = 0.1$, and a seed set $I_0 = 1$. Unless otherwise stated, the results for each scenario are averaged over 50 realizations. Furthermore, for the optimization campaigns if the budget B is not fully utilized by the campaign, we randomly choose T nodes and distribute the remaining budget equally among them. In all figures presented, the error bars correspond to 95% confidence intervals.

6.3.1 Control Performance on Different Network Structures

We begin with the equal-budget scenario, where both positive and negative campaigns have the same budget , i.e., $\mu^- = \mu^+$, to evaluate the efficacy of control strategies and to understand the propagation dynamics of anti-vaccine diffusion across the considered networks. The campaigns considered are DynRandT and DynAdvLocT, introduced in Chapter 3; the individual-based campaign Min_NK

introduced in Chapter 4; the cluster-based campaign Cl.$_{wt}$ introduced in Chapter 5; and the newly introduced weight-augmented individual campaign Min_NK$_{wt}$.

We begin by applying the DynAdvLocT campaign to random and scale-free networks, using the same approach previously applied to the small-world network in Chapter 3. This campaign relies on two main parameters, ζ and Z, which represent the target number of anti-vaccine neighbors and the target number of neutral neighbors, respectively. Our goal here is to identify the optimal values for these parameters in these networks. In this campaign, ζ governs the targeting of nodes adjacent to anti-vaccine neighbors to block the diffusion of negative influence by focusing on agents vulnerable to negative exposure, while Z governs the selection among these nodes such that increasing the value of Z prioritizes nodes with more neutral neighbors to facilitate the spread of positive influence (see Section 3.3.2.2 for more details). Figures 6.1(A) and 6.1(B) show the epidemic size obtained for varying values of Z and ζ in random and scale-free networks, respectively. For both network types, the results show that as Z increases, the epidemic size decreases, while the best results for ζ are observed when $\zeta \geq 10$, equivalent to the average degree, beyond this point the reduction levels remain relatively constant. In the random network (Fig. 6.1 (A)), increasing Z beyond the average degree (i.e., $Z > 10$) leads to better performance. In contrast, in the scale-free (Fig. 6.1 (B)), setting Z to significantly higher values (e.g., $Z \geq 100$) results in a greater reduction in epidemic size.

In heterogeneous scale-free networks, nodes with high degree typically have more neutral neighbors than others, specifically in early intervention scenarios before negative influence spreads widely in the network. In this campaign, increasing the parameter Z scores agents in a way that implicitly prioritizes high-degree nodes by selecting those with more neutral neighbors. Consequently, higher values of Z result in more effective control and a smaller epidemic size. This suggests that targeting nodes with high degrees is an effective strategy in heterogeneous structures, even in coupled dynamics, aligning with findings from existing literature (Chen et al., 2009; Masuda, 2015).

Next, we apply the intervention campaigns to all networks under consideration to systematically evaluate the impact of network topology on the intervention effectiveness at distinct intervention timings. Figure 6.2 presents the resulting final epidemic size, where Fig. 6.2 (A) corresponds to interventions applied at $t_s^+ = 100$, Fig.6.2 (B) at $t_s^+ = 175$, and Fig.6.2 (C) at $t_s^+ = 200$. Due to the observed variations in diffusion speed across networks, we found that the number of anti-vaccine opinion adopters at time t_s^+ is similar for all networks, except for the scale-free network, where it is higher. For example, at $t_s^+ = 200$, the number of adopters is approximately (334) for all networks, except in the scale-free network, where it is (1333). Therefore, we applied the interventions slightly earlier in the scale-free network to achieve a comparable number of opinion adopters for comparison purposes. The average

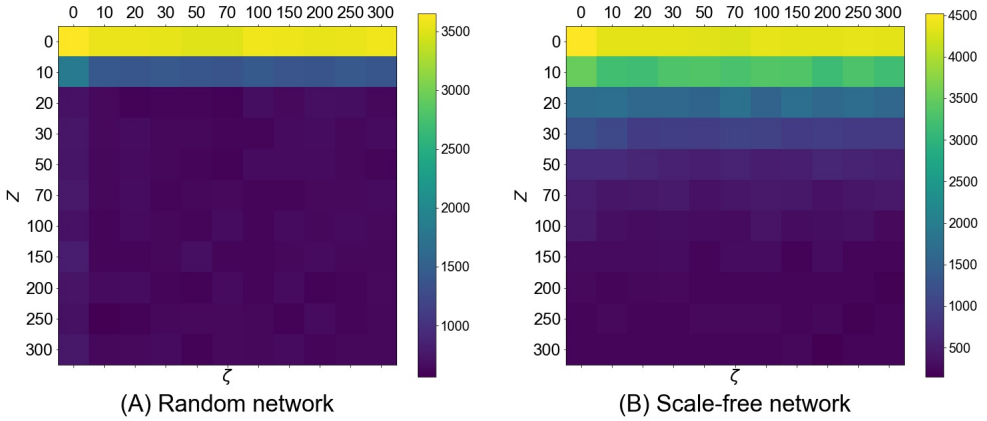

FIGURE 6.1: Average epidemic size using the DynAdvLocT campaign in (A) random networks and (B) scale-free networks. The epidemic size is shown as a function of the target number of anti-vaccine neighbors ζ and the target number of neutral neighbors Z. Applying the positive campaigns at time $t_s^+ = 100$. Other settings are $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, $t_r = 1$, $T = 5$, and the results are the average of 100 realizations.

number of anti-vaccine opinion adopters at each time point is provided in the figure caption.

As shown in Fig. 6.2, the underlying network structure significantly influences the diffusion dynamics, consistent with previous findings in the literature (Pérez-Ortiz et al., 2022), affecting the effectiveness of control strategies and the resulting epidemic size. An interesting observation is that minimizing negative opinion adopters through the Min_NK campaign (see black bars of Fig. 6.2) is more effective in scale-free and random network than in small-world networks, where it was previously demonstrated to be less effective. For example, at the early intervention shown in Fig. 6.2 (A), the Min_NK campaign performs significantly better in scale-free networks compared to its performance in the small-world network. This outcome can be attributed to the targeting scheme of the campaign, as detailed in chapter 4, which targets nodes with a high number of anti-vaccine neighbors. In scale-free networks, such nodes are more likely to be high-degree (hub) nodes, as they are most vulnerable to negative influence due to their large number of connections. At early stages, these hub nodes still have many neutral neighbors who can be positively influenced by targeted agents. As a result, this leads to more effective blocking of negative influence and promotion of positive influence, ultimately yielding a significant reduction in epidemic size. This further reinforces our earlier observation, which suggests that targeting high-degree nodes is a more effective strategy in heterogeneous structures.

We observe that the weight-augmented Min_NK$_{wt}$ campaign exhibits similar performance to the unweighted campaign Min_NK in scale-free networks (see red and black bars in Fig. 6.2), unlike its positive effect in other networks. This is because hub nodes in such networks, due to their high connections, are the most suitable

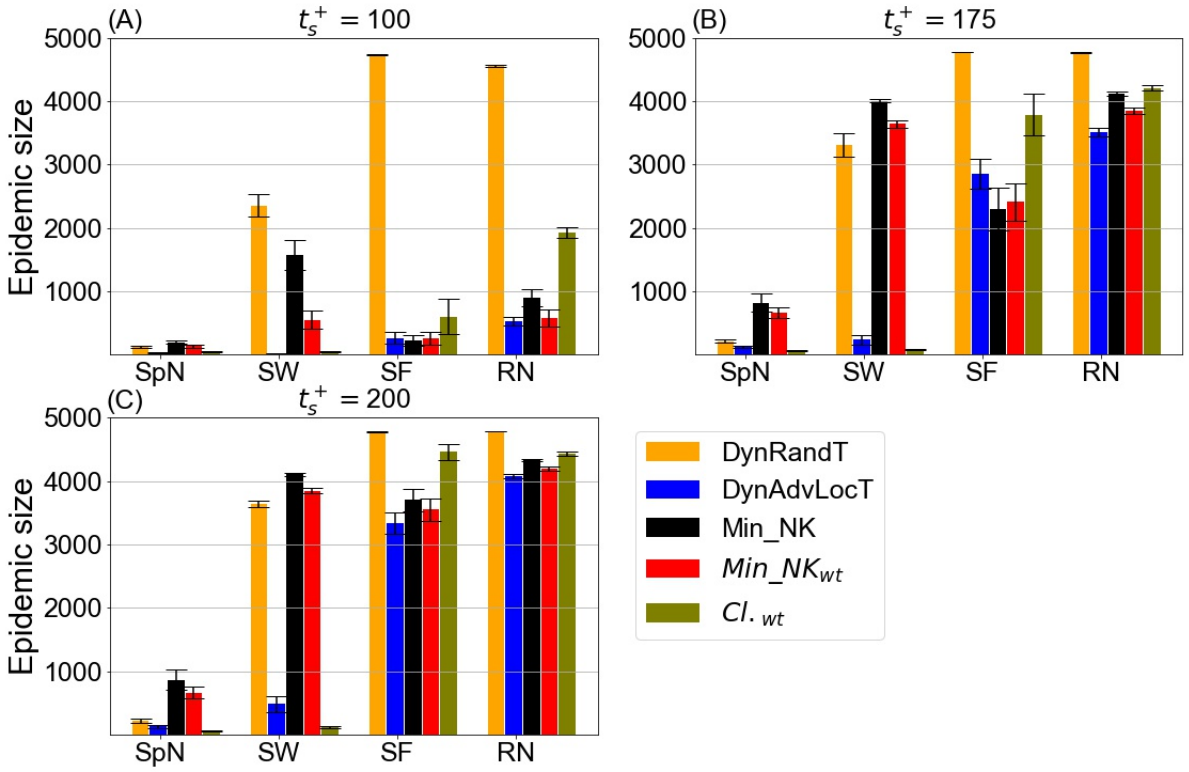

FIGURE 6.2: Comparison of final epidemic size across different network topologies and varying intervention times t_s^+ for the different intervention campaigns. The campaigns were applied at (A) $t_s^+ = 100$ for all networks, with the average number of existing anti-vaccine opinion adopters being around 34 ± 2 (approximately 1% of the population); (B) $t_s^+ = 175$ for all networks except scale-free at $t_s^+ = 145$, with around 208 ± 10 existing anti-vaccine opinion adopters (approximately 4% of the population); and (C) $t_s^+ = 200$ for all networks except scale-free at $t_s^+ = 160$, with around 315 ± 14 existing anti-vaccine opinion adopters (approximately 10% of the population). Other settings are $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, $t_r = 1$, $T = 5$. For DynAdvLocT campaign $\zeta = 10$ and $Z = 100$. SpN, SW, SF, and RN denote regular lattice, small-world, scale-free, and random networks, respectively.

candidates for both the Min_NK and Min_NK$_{wt}$ campaigns. They are more likely to have a large number of anti-vaccine neighbors, making them appropriate targets for Min_NK, as well as many neutral neighbors, making them equally suitable for Min_NK$_{wt}$.

In addition, in scale-free networks —where most nodes are peripheral with few connections—the dynamics may differ. The motivation behind considering adjacent neutrals is to facilitate the widespread diffusion of positive influence and initiate positive contagion within the network. However, peripheral nodes contribute little to spreading positive influence throughout the network due to their limited reach. Therefore, even if the campaign targets these peripheral nodes at certain stages, it is

unlikely to yield further improvements in control. Consequently, this campaign yields no significant improvement in scale-free networks (see black and red bars in the 'SF' category of Figs.6.2(A)-(C)). This also indicates that initiating positive contagion throughout the network is more challenging in such structures if we fail to control the hub nodes. In contrast, this campaign demonstrates better performance in random networks, owing to its ability to spread positive influence across the entire network. Moreover, the DynAdvLocT campaign exhibits similar performance to the Min_NK and Min_NK$_{wt}$ campaigns in both networks (RN and SF), with the Min_NK campaign performing marginally better in some scenarios in SF (see blue and black bars in the 'SF' category of Fig. 6.2 (B)).

Moreover, the cluster-based strategy proves to be less effective in scale-free and random networks the longer the intervention is delayed (see olive bars in the 'SF' and 'RN' categories of Figs. 6.2 (A)-(C)). This observation highlights the differences in the structure of emerging anti-vaccine communities across these network typologies. It suggests that cluster-based strategies that are effective in small-world and spatial networks (see olive bars in the 'SpN' and 'SW' categories of Figs. 6.2 (A)-(C)) may underperform in heterogeneous networks with hub nodes and in random networks. In scale-free networks, the evolution of negative diffusion in such topologies does not primarily depend on interactions between different clusters—where protecting the critical region helps control the growth of these communities—but rather on controlling hub nodes, as transitioning these nodes to the negative opinion can rapidly lead to large communities. Similarly, the cluster-based approach is inefficient in random networks due to the inherent randomness of connections which impacts the shape of the emergent anti-vaccine clusters. This also reflects the complexity of identifying critical regions in such topologies.

In small-world networks, characterized by high local clustering, cluster-based campaigns provide effective control, as demonstrated previously (see the olive bars in the 'SW' category of Figs. 6.2 (A)-(C)). Although the optimized Min_NK$_{wt}$ campaign (red bars) outperforms Min_NK (black bars) in this network, the cluster-based approach remains the optimal intervention. Despite the average degree of the network being the same as in random and scale-free networks, the structure of the emerging communities during influence propagation helps contain the diffusion by shielding critical regions within the network. Similar performance trends are observed in the regular lattice networks (see the 'SpN' category of Figs. 6.2 (A)-(C)), which have similar properties to small-world networks, where the cluster-based campaign yields the best results. Additionally, regular lattice networks exhibit greater resilience to the spread of negative influence. The dynamics of this influence can be efficiently contained even with late interventions, leading to a significant reduction in epidemic size with almost all campaigns.

In conclusion, our campaigns demonstrate effective control across all network types, as long as the intervention is applied early enough, e.g., when approximately 1% of the population has adopted an anti-vaccine opinion at $t_s^+ = 100$. In practice, given that geographic regions exhibit varying levels of vaccine uptake and hesitancy (Thompson et al., 2023; Matas et al., 2023), early interventions could be feasible in areas with lower vaccination coverage or where negative information has not yet spread widely. However, we find that control becomes challenging with a longer delay in the intervention, e.g., after approximately 10% of the population has adopted an anti-vaccine opinion (Fig. 6.2 (C)), particularly in random and scale-free networks. Nonetheless, in this scenario, the spread can still be effectively contained in small-world and lattice networks through the cluster-based campaign. As we will see below, this challenge stems from the shape and clustering patterns of pre-existing anti-vaccine communities that form prior to the intervention in different topologies. This emphasizes the importance of understanding how these communities evolve in each network. In the next section, we analyze and visualize the existing anti-vaccine communities in each network at different times before the intervention to gain a deeper understanding of the distinct dynamics in each network. We demonstrate how the structure and distribution of these communities complicate containment efforts, making effective control more difficult, particularly with a limited budget.

6.3.2 Analysis of Anti-vaccine Communities

Here, we analyze the distribution of negative opinion adopters across different topologies before the intervention to understand their clustering behaviour. We find that under identical conditions for negative influence propagation (i.e., the same network size, average degree, and strength of negative diffusion), with the absence of any positive intervention, the resulting spread patterns and emerging clusters exhibit significant variation. Specifically, in small-world and lattice networks, anti-vaccine opinion adopters are scattered across many small communities. In contrast, in scale-free and random networks, the same number of adopters tend to form one dominant large cluster, along with several smaller, distributed clusters.

Figure 6.3 illustrates an example visualization of entire networks before the positive campaign is applied. The left side of the figure displays the whole network, highlighting both neutral and anti-vaccine individuals, while the right side depicts the connectivity patterns of the anti-vaccine communities. As observed in this example, anti-vaccine communities in the scale-free network coalesce into a single giant cluster, while the random network exhibits one large cluster alongside several smaller ones. In contrast, in small-world network, the adopters are distributed across many smaller communities.

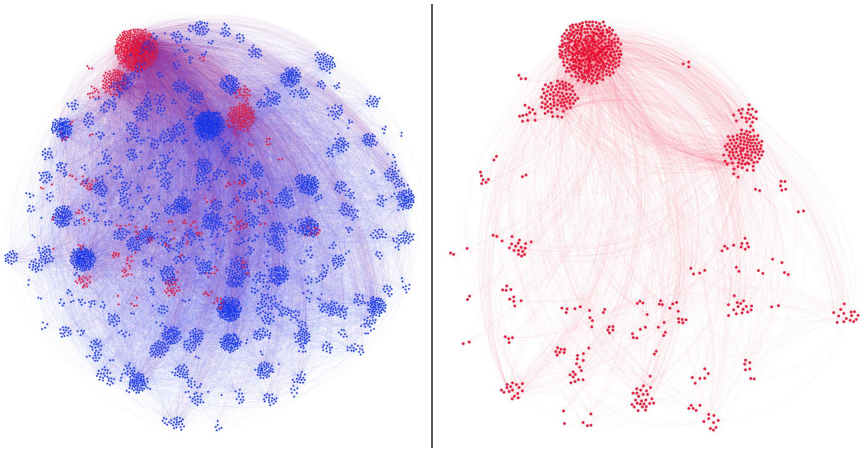

(A) Scale-free network with one dominant anti-vaccine community.

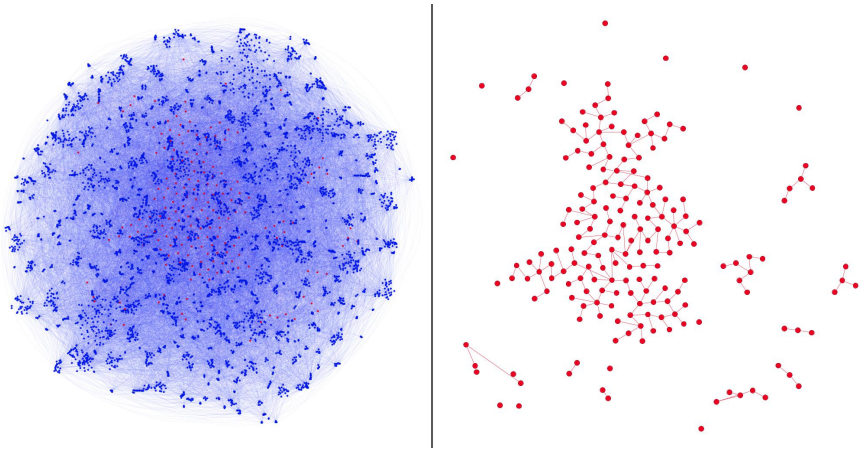

(B) Random network with one large anti-vaccine community alongside multiple smaller communities.

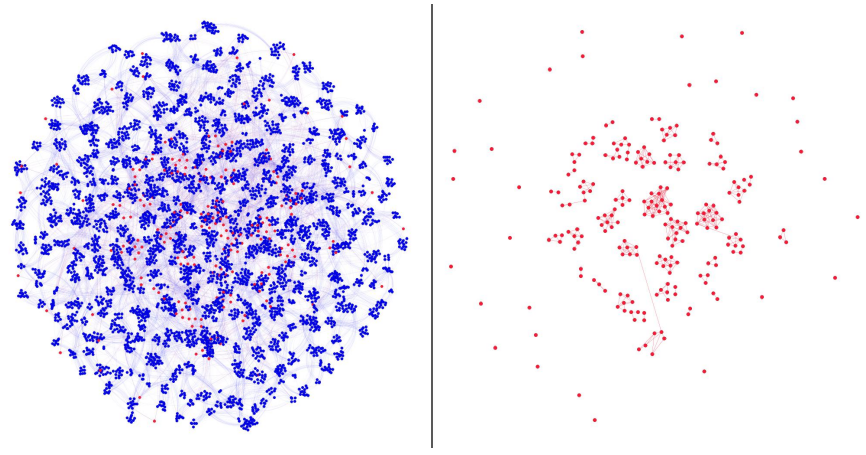

(C) Small-world network with multiple disconnected anti-vaccine communities.

FIGURE 6.3: An example visualization showing anti-vaccine communities across different networks prior to intervention campaign implementation at $t_s^+ = 145$ in SF, and $t_s^+ = 175$ in others. The entire network is shown in the left panels, while the anti-vaccine communities are displayed in the right panels. Blue and red represent neutral and anti-vaccine agents, respectively.

This explains the poor performance observed with late intervention for all campaigns in scale-free and random networks compared to others. With a limited budget, it becomes challenging to curb the growth of large clusters or control other smaller emerging ones. In scale-free networks, hub nodes play the dominant role; once negatively affected, they can rapidly propagate their influence throughout the network, facilitating the formation of large anti-vaccine clusters over a shorter time compared to other network types. In such networks, identifying critical regions between established anti-vaccine clusters is challenging, as emerging anti-vaccine agents tend to form a single large community from the early stages. Similarly, in random networks with arbitrary global connections, containing the diffusion is difficult, as the influence can easily form large, distributed communities. This also clarifies why the cluster-based approach is less effective in these topologies. Many susceptible agents are more likely to be adjacent to the largest clusters, which, in most cases over time, form a single large cluster in these networks. Therefore, given that this campaign targets critical agents — those connected to multiple anti-vaccine groups — such agents are rarely found as most nodes are often adjacent to a single giant cluster, making the campaign ineffective in preventing the growth of these large communities.

To generalize our findings, figure 6.4 presents the average size of the largest clusters before the initiation of positive campaigns, along with the second-largest clusters across different realizations. As shown in the figure, in both random and scale-free networks, negative opinion adopters are concentrated in large clusters with an average size of 132 ± 24 and 149 ± 34, respectively, as illustrated by the olive and brown bars in the "Largest" category of Fig. 6.4. In contrast, in small-world and spatial networks, the largest clusters tend to be significantly smaller, with sizes below 30, as illustrated by the dark orange and dark green bars in the "Largest" category of Figure 6.4.

Examining the second-largest cluster in Fig. 6.4, we observe that in random networks, its average size is approximately 13 ± 4, while in scale-free networks, the clusters are much smaller, averaging around 2 ± 0.5, indicating that only a few anti-vaccine opinion adopters are scattered across the network. This explains why controlling negative diffusion in a random graph is the most challenging (see the epidemic size at $t_s^+ = 175$ and $t_s^+ = 200$ for the random network (RN) in Figure 6.2). The presence of a large cluster coupled with several smaller communities in the random network makes negative containment more difficult. In small-world and spatial networks, the second-largest cluster is only slightly smaller than the largest one, reflecting the overall small cluster sizes in these networks. In later intervention scenarios, the formation of anti-vaccine communities exhibits the same pattern but with larger sizes, making the control of these large clusters increasingly difficult (see Figure B.3 in Appendix B for additional cluster analysis related to later interventions, where a higher number of anti-vaccine adopters are exist).

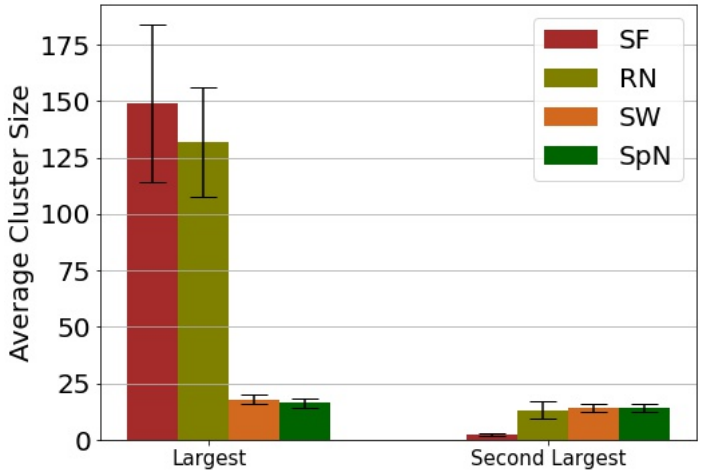

FIGURE 6.4: The average size of the largest and second-largest clusters of anti-vaccine communities before applying the positive campaigns at time $t_s^+ = 175$ for all networks, except for scale-free network at $t_s^+ = 145$. At the time of applying the positive campaign, around 4% of the population had adopted the anti-vaccine opinion. A total of 20 simulations were conducted. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

Does this clustering pattern persist from earlier stages? Figure 6.5 presents the same analysis for the scale-free and random networks at an earlier stage, with around 97 ± 5 anti-vaccine opinion adopters, to investigate whether this clustering pattern forms early in the process. The results reveal that the same pattern emerges from the onset of community formation. Specifically, in the random network, a large cluster coexists with many smaller ones, whereas in the scale-free network, a single large cluster dominates.

Until now the analysis has focused on an equal-budget scenario to understand the dynamics of influence propagation and the formation of anti-vaccine communities. We have highlighted the most effective intervention strategies for each network and explained the causes of performance degradation. We found that the main challenge lies in the formation of large dominant clusters pre-intervention, which make further containment very challenging. Below, we continue by investigating the effects of the amount of available campaigning resource on campaign efficiency.

6.3.3 Dependence of Campaign Efficiency on Budget Size

Here, we evaluate the performance dependency on the allocated positive budget to determine the resources needed to contain negative diffusion and the extent of the resulting reduction in epidemic size against a fixed negative budget. Figure 6.6 illustrates the epidemic size as a function of the allocated budget, μ^+, across the

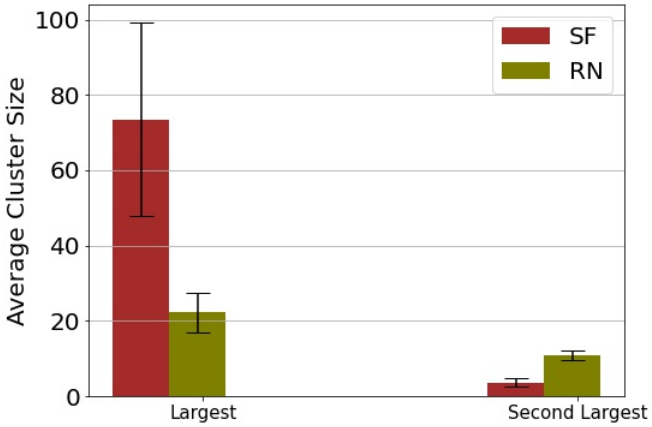

FIGURE 6.5: The average size of the largest and second-largest clusters of anti-vaccine communities before applying the positive campaigns at time $t_s^+ = 145$ for random network, for scale-free network at $t_s^+ = 120$. At the time of applying the positive campaign, around 2% of the population had adopted the anti-vaccine opinion. A total of 20 simulations were conducted. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

networks and the intervention strategies. It also shows the final number of anti-vaccine $\langle n^- \rangle$ and pro-vaccine $\langle n^+ \rangle$ opinion adopters. We consider in this figure the scenario where the intervention occurs at $t_s^+ = 175$, with around 4% of the population adopting anti-vaccine opinion. In addition, since the DynAdvLocT campaign depends not only on the budget but also on the target size, we include a larger target set size, for this campaign.

For the random network RN (Figs. 6.6 (A)-(C)), the weight-augmented campaign Min_NK$_{wt}$, represented by the red color, delivers the best performance across all budget scenarios, even when operating under a budget disadvantage, i.e., $\mu^+ < \mu^-$. We find that the performance in reducing epidemic size improves steadily with an increasing budget (red color in Fig. 6.6 (A)). With a positive budget 50% larger than the negative budget, we achieve a 23% reduction compared to the equal-budget scenario, i.e., $\mu^+ = \mu^-$. When the positive budget doubles the negative budget, up to 60% of the population is saved. Furthermore, with a threefold budget, up to 80% of the population can be saved, reducing the final epidemic size to around 972 $\pm$133 (see the red curve at $\mu^+ = 0.003$ in Fig. 6.6 (A)). The DynAdvLocT campaign with a larger target set achieves identical performance to Min_NK$_{wt}$ (see the light blue curve in Fig. 6.6 (A)).

In the scale-free network (Figs. 6.6 (D)-(F)), control performance generally surpasses that of the random network (Figs. 6.6 (A)-(C)) using the Min_NK$_{wt}$ and Min_NK campaigns, even with a lower budget. The performance improves steadily as the budget increases. With a 50% increase in the budget (at $\mu^+ = 0.0015$), the reduction

improves by 23% compared to the equal-budget scenario. With double the resources, approximately 70% of the population can be saved (black curve in Fig. 6.6 (D)).

In the small-world network (Figs. 6.6 (G)-(I)), as demonstrated in the previous chapter, the cluster-based campaign is the optimal strategy. It achieves effective containment with a significantly lower budget than required in the counterpart scale-free and random networks. For instance, in a budget disadvantage scenario, i.e., $\mu^+ = 0.007$, with the cluster-based approach only around 14% of the population get infected (see the olive curve at $\mu^+ = 0.0007$ in Fig. 6.6(G)). Despite the final number of anti-vaccine opinion adopters being high, directing resources to prevent the clustering and the formation of a giant anti-vaccine cluster effectively shields critical regions in the network and results in a smaller epidemic size through this campaign. In this scenario, anti-vaccine opinion adopters are distributed across 39 ± 3 communities. The spatial network in Figs. 6.6 (J)-(L), characterized by the absence of global connections and links restricted to nearest neighbors, shows even more promising results. In this case, the cluster-based campaign achieves a greater reduction, saving more than 90% of the population compared to the worst-performing campaign, Min_NK, with only a quarter of the budget compared to the negative budget, i.e., $\mu^+ = 0.00025$ (see the black and olive curves at $\mu^+ = 0.00025$ in Fig. 6.6(J)).

These findings demonstrate that the resources required to contain negative diffusion and disease spread increase significantly in random and scale-free networks as the intervention is delayed, with the random network requiring even more resources (see epidemic size in Figs. 6.6(A) and (D) for random and scale-free networks, receptively)[1]. In contrast, highly clustered networks (small-world and lattice networks) enable better containment with comparatively limited resources (see epidemic size in Figs. 6.6(G) and (J) for small-world and lattice networks, respectively).

Figure 6.7 presents the results of intervention scenarios implemented at earlier stages in scale-free and random networks, examining the effects of earlier intervention in these two networks. As shown, a 50% increase in the positive budget results in a substantial reduction in the random network (red curve at $\mu^+ = 0.0015$ in Fig.6.7(A)), whereas the reduction in the scale-free network is more marginal (red curve at $\mu^+ = 0.0015$ in Fig.6.7(D)). A threefold increase in the budget in the random network leads to a nearly disease-free state with epidemic size 50 ± 26 (see the red curves at $\mu^+ = 0.003$ in Fig. 6.7(A)), while the epidemic size in the scale-free network is reduced to a minimal level compared to the later intervention in Fig. 6.6(D) (red curves at $\mu^+ = 0.003$ in Fig. 6.7(D)). Further increases in the budget in the scale-free network have a limited effect, as the epidemic size is already small.

[1] A scenario depicting a later intervention at $t_s^+ = 200$ is provided in Fig. B.2 in Appendix B.

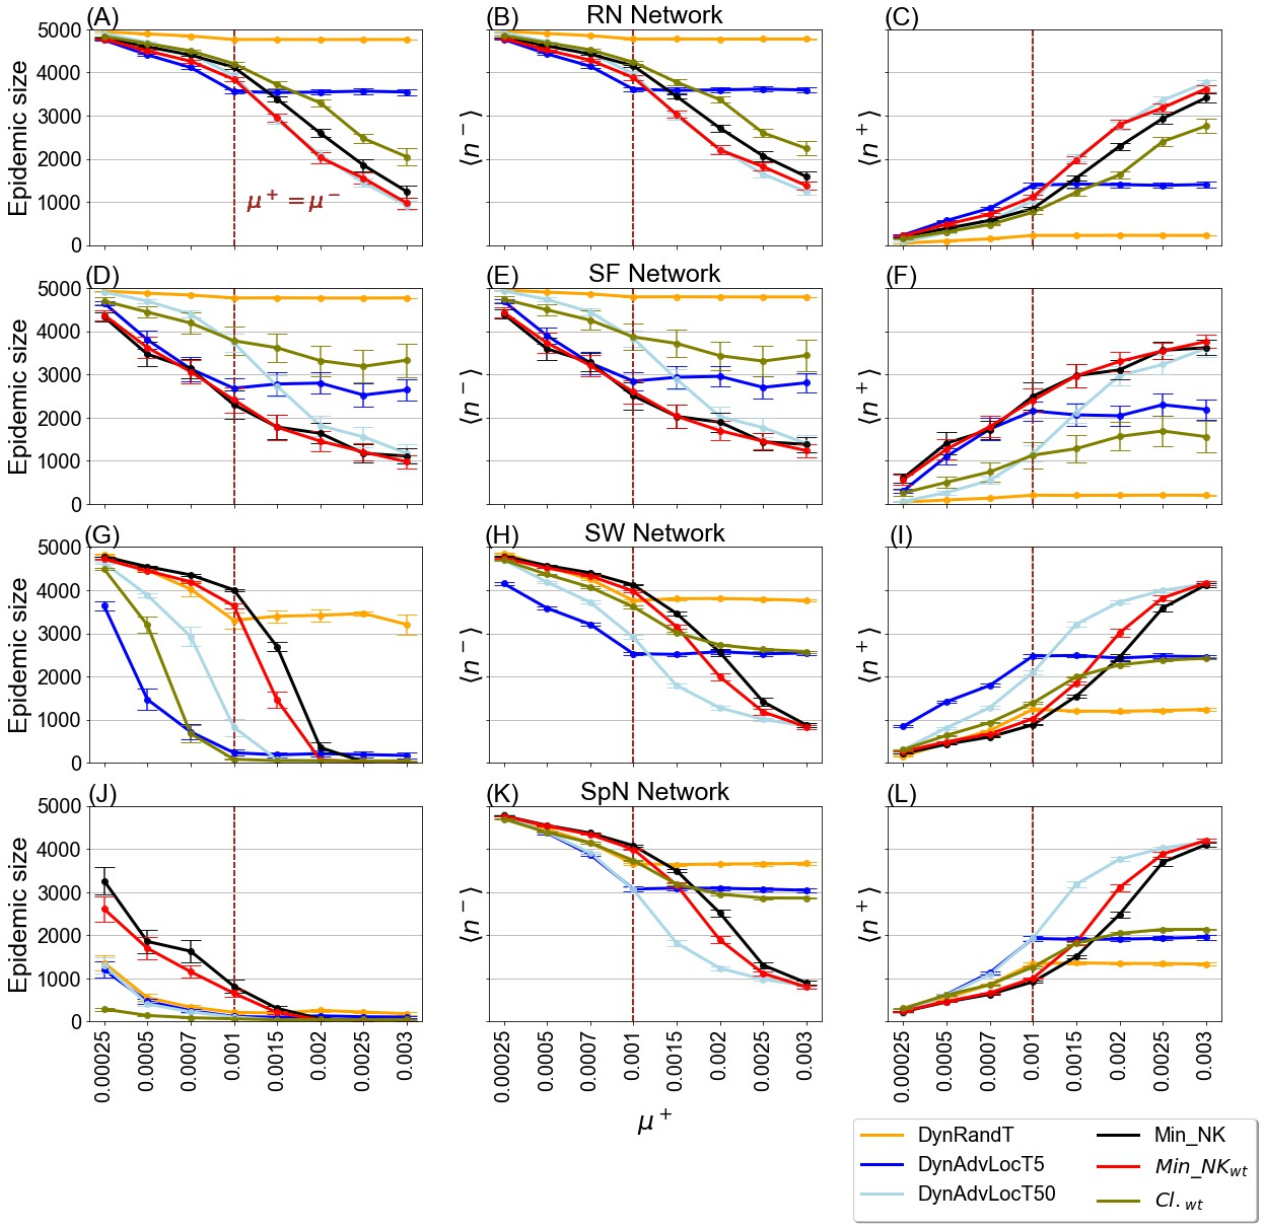

FIGURE 6.6: Comparison of the control performance of the intervention campaigns across different network structures as a function of μ^+. Applying campaigns at $t_s^+ = 175$ in Random (RN), small-world (SW) and spatial (SpN) networks, and at $t_s^+ = 145$ in scale-free network (SF), with an average of 208 ± 10 existing anti-vaccine opinion adopters. The left panels illustrate the resulting epidemic size, the middle panels show the average number of anti-vaccine opinion adopters $\langle n^- \rangle$, and the right panels illustrate the average number of pro-vaccine opinion adopters (vaccinated) $\langle n^+ \rangle$. With $\mu^- = 0.001$, $\omega = 0.006$, $T = 5$, $t_r = 1$, for DynAdvLocT campaign $\zeta = 1, Z = 10$ for SW and SpN, and $\zeta = 10, Z = 100$ for RN and SF.

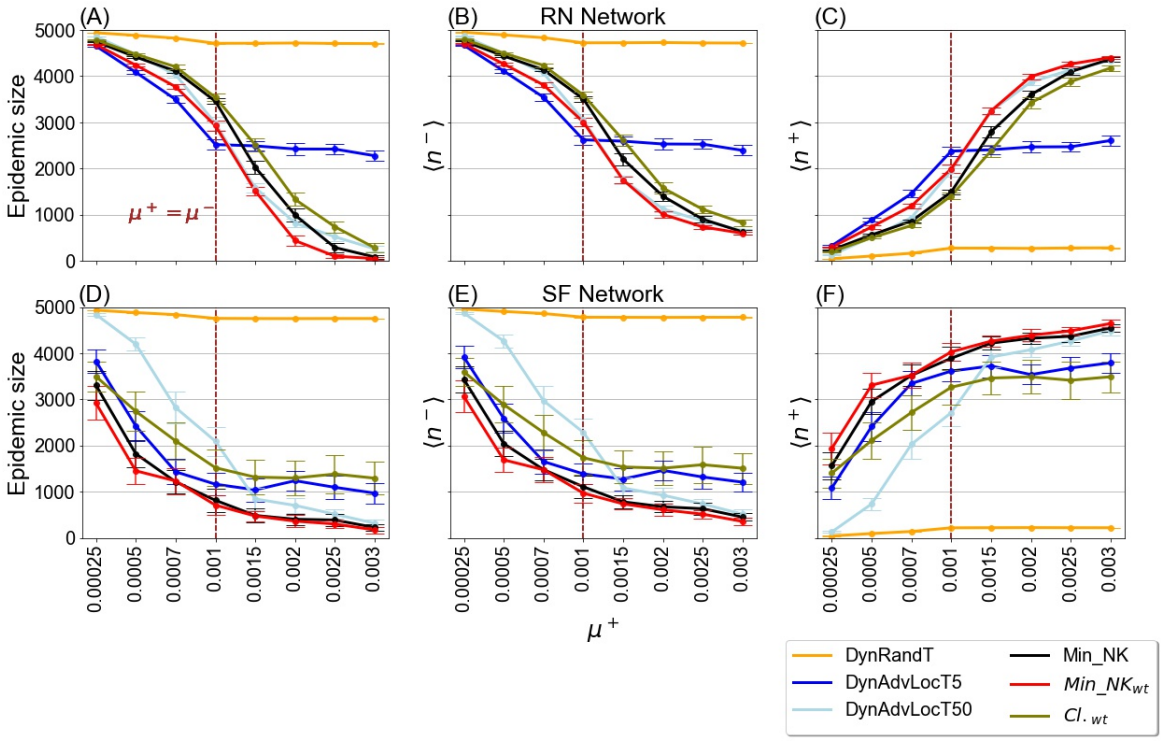

FIGURE 6.7: Comparison of the control performance of the intervention campaigns across different network structures as a function of μ^+. Applying campaigns at $t_s^+ = 145$ for the random network and $t_s^+ = 120$ for the scale-free network, with an average of 97 ± 5 existing anti-vaccine opinion adopters. The left panels illustrate the resulting epidemic size, the middle panels show the average number of anti-vaccine opinion adopters $\langle n^- \rangle$, and the right panels illustrate the average number of pro-vaccine opinion adopters (vaccinated) $\langle n^+ \rangle$. With $\mu^- = 0.001$, $\omega = 0.006$, $T = 5$, $t_r = 1$, for DynAdvLocT campaign $\zeta = 10, Z = 100$.

6.3.4 Control Performance in Real-world Networks

Real-world networks exhibit a wide range of properties, such as degree distributions, clustering, and density, with link correlations that may not be accurately captured in synthetic network models. Here, we consider two real-world networks representing friendship networks extracted from the Facebook social network platform, obtained from the Network Repository (Rossi and Ahmed, 2015). These networks capture realistic social interactions and differ in size and connectivity from the synthetic networks discussed earlier. Although they represent online ties that may not accurately reflect offline relationships, they can serve as proxies for offline social interactions given the observed overlap between online and offline social networks (Reich et al., 2012). The properties of these networks are summarized in Table 6.1. In this section, we apply the intervention campaigns to these networks to assess whether the control achieved is consistent with the insights gained from the dynamics of synthetic networks in the previous sections. From the degree distribution of these networks (Rossi and Ahmed, 2015), we observe that they exhibit an approximate

power-law pattern, implying a large degree of degree heterogeneity [2]. We apply our campaigns to these two selected networks at different intervention times. In addition, since the epidemic size decreases monotonically with an increase in the budget, we consider the equal-budget scenario in this section.

Figure 6.8 shows the resulting epidemic size for the intervention applied at different timings in these two networks. Since we find that campaigns' performance declines significantly on these two networks at $t^+ = 100$, where most of the population has been infected, we exclude analyses beyond this intervention time for these networks and instead explore earlier intervention at $t^+ = 50$. We make the following observations. First, the Min_NK$_{wt}$ strategy is the most effective approache for controlling negative diffusion and subsequently reducing epidemic size in these networks, particularly with earlier intervention at $t_s^+ = 50$ (gray bar in Figs. 6.8(A) and (C)). The DynAdvLocT and Min_NK campaigns exhibit similar performance, with slightly higher epidemic size than the Min_NK$_{wt}$ campaign. Second, this pattern is consistent with the observed effectiveness of these strategies in managing diffusion within scale-free networks (see epidemic size in the 'SF' category of Figs. 6.2(A)-(C)). Third, in these dense and heterogeneous networks, the spread of negative influence occurs rapidly, making it increasingly challenging to control as intervention is delayed. This observation also aligns with the dynamics of scale-free networks discussed in the previous section. Due to the higher degree and density, the scenario is more severe in this context. For instance, intervening at $t_s^+ = 100$ — previously one of the effective intervention points in less dense scale-free networks — proves insufficient here, as negative influence has already reached a substantial portion of the population by this time. Specifically, by $t_s^+ = 100$, approximately 73% of the population in the socfb-JohnsHopkins55 network and 62% in the socfb-Maine59 network have been negatively influenced.

Analyzing the distribution of negative adopters reveals that, by $t_s^+ = 100$, a substantial portion of the population has been negatively influenced and forms a giant cluster with an average size of 3819 ± 201 and 5766 ± 318 in socfb-JohnsHopkins and socfb-Maine59, respectively (see the "Largest" category bars in Fig. 6.9 (B)). In contrast, at $t_s^+ = 50$, these clusters are significantly smaller (see the "Largest" category bars in Fig. 6.9 (A)), which allows for more effective control if intervention occurs at this earlier time point. Furthermore, due to the lower density of the socfb-Maine59 network, early intervention results in a more substantial reduction in epidemic size. For example, using Min_NK campaign achieves approximately a 79.8% reduction at $t_s^+ = 50$ compared to intervention at $t_s^+ = 100$ in the socfb-Maine59 network. In contrast, the socfb-JohnsHopkins network, despite being smaller in size, achieves about a 52.2% reduction at $t_s^+ = 50$ relative to $t_s^+ = 100$ using the same campaign.

[2] See the degree distributions of these networks in Appendix B.1.

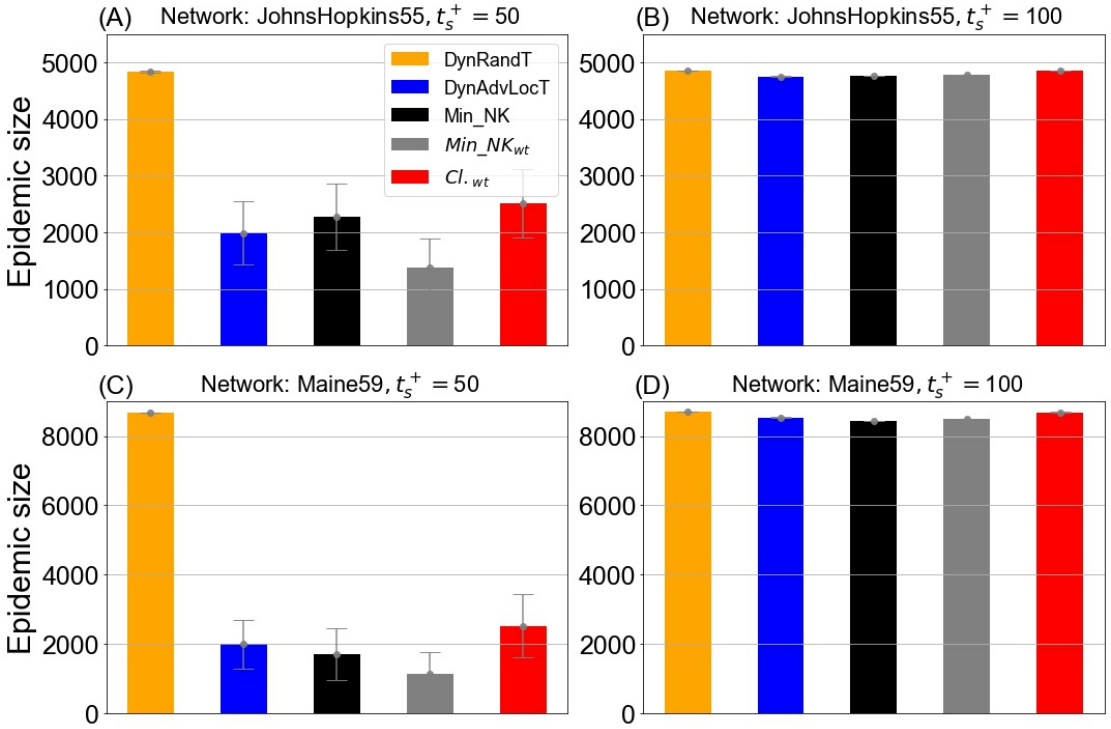

FIGURE 6.8: Average epidemic size for intervention campaigns in the JohnsHopkins55 and Maine59 networks. JohnsHopkins55 network: intervention starting at time (A) $t_s^+ = 50$ with 36 ± 21 anti-vaccine opinion adopters, and (B) $t_s^+ = 100$ with 3819 ± 205 anti-vaccine opinion adopters.Maine59 network: intervention starting at time (C) $t_s^+ = 50$ with 24 ± 5 anti-vaccine opinion adopters, and (D) $t_s^+ = 100$ with 5767 ± 324 anti-vaccine opinion adopters. With $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$. For DynAdvLocT campaign $\zeta = 10, Z = 300$. Note that the network size for JohnsHopkins55 is $5.2K$, and for Maine59 is $9.1K$.

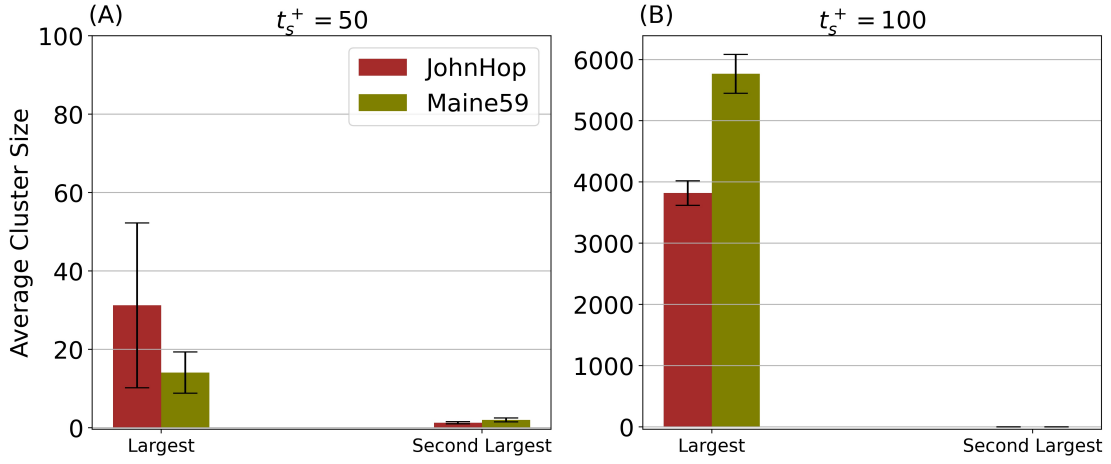

FIGURE 6.9: The average size of the largest and second-largest clusters of anti-vaccine communities before applying the positive campaigns at time (A) $t_s^+ = 50$ and (B) $t_s^+ = 100$ for the JohnsHopkins55 and the Maine59 networks. Results are the average of 50 realization. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$. Note that the y-axis range is different for t=50 and t=100.

6.4 Summary

In this chapter, we have studied the role of the network structure in the effectiveness of the proposed intervention campaigns. We have systematically analyzed the influence of network topology on the efficacy of campaigns across a variety of synthetic and real-world networks, evaluating their robustness, identifying the most effective strategies for each network structure, and highlighting the limitations of each approach. We have also extended our framework and developed a new intervention campaign that aims to minimize the number of anti-vaccine opinion adopters while also potentially fostering the spread of positive influence within the network.

Our results reveal that network structure significantly affects both mitigation efforts and the level of control achieved, and there is no single campaign that works effectively for all topologies. We demonstrate that across all network structures, early interventions targeting individuals who are susceptible to negative influence but simultaneously possess the potential to positively influence their social network is significantly more effective than solely protecting the most vulnerable individuals from negative influences. This can be delivered through DynAdvLocT and Min_NK$_{wt}$ campaigns, which both exhibit efficient control performance. However, their effectiveness diminishes the longer the intervention is delayed. In scenarios with late intervention and a greater dominance of negative influence, we find that shielding critical network areas through the cluster-based campaign is effective in highly clustered networks (small-world and spatial networks) but less effective in scale-free and random networks. Identifying these critical regions is challenging in scale-free and random networks due to the evolving nature of how these communities form within such structures.

We have further analyzed the clustering and propagation behavior of anti-vaccine opinion adopters across these topologies. We find that the formation of anti-vaccine communities varies in both shape and speed, with distinct clustering patterns observed. In scale-free networks, due to their connectivity pattern and the presence of hub nodes that accelerate the diffusion of influence, anti-vaccine opinion adopters tend to form a single community. This leads to the early emergence of a giant anti-vaccine cluster, which occurs more rapidly than in small-world and lattice networks. Similarly, in random networks, we observed the rapid formation of a large cluster coupled with several smaller clusters. These dynamics complicate the control and demands greater resources for effective containment.

We then examine the impact of the allocated budget across different network topologies. We find that the longer the intervention is delayed, the more challenging and resource-intensive it becomes to control negative opinion diffusion. However, despite this difficulty, it can still be contained more efficiently and with fewer

resources in small-world and spatial networks than in scale-free and random networks.

This chapter draws key observations and insights into the patterns of propagation across various network topologies and highlights critical aspects of control performance. We found that several key factors play a crucial role in the effectiveness of control performance: the network structure; the timing of interventions; the extent of the negative dominance; and the allocation of the campaign's positive budget. Our findings demonstrate that for a given budget, the impact of campaigns can be improved by adopting targeting strategies that consider these key factors.

Although our model provides valuable insights, a key limitation is the use of a single network to represent both opinion diffusion and disease spread, as it considers only face-to-face social interactions. This simplified structure allows us to systematically explore how anti-vaccine influence spreads across different network topologies, how anti-vaccine clusters evolve within these topologies, and how these variations impact the performance of intervention strategies at different intervention timings. However, it is important to explore more realistic scenarios that incorporate diverse types of social interactions. In the next chapter, we address this limitation by extending our model to a multi-layer network framework, in which information dissemination and disease transmission are represented by distinct networks. Building on the understanding gained thus far, we aim to investigate system dynamics within this complex framework and evaluate the effectiveness of the interventions.

Chapter 7

Intervention Campaigns in Multi-Layer Networks

7.1 Introduction

In the previous chapters, the focus has been on face-to-face interactions, overlooking the diversity of social networks present in today's complex social systems. In reality, people engage in different types of social relationships and interact through various social channels, including social media. Therefore, in this chapter, we explore a more realistic scenario using a multi-layer network model where information diffusion and disease transmission occur in separate layers. A deeper understanding of these dynamics is vital for assessing the effectiveness of intervention campaigns in curbing information diffusion across different social networks and their impact on the emergence of unvaccinated communities in the physical network where disease transmission occurs.

Multiplex networks are a form of multi-layer networks in which the same set of nodes is connected through multiple types of relationships, each represented as a distinct layer (Aleta and Moreno, 2019; Ma and Parilina, 2024). These layers capture different communication patterns or interactions between nodes, enabling the analysis of various types of connections between agents and the impact of their interplay in coupled dynamics. In coupled information-disease dynamics, multiplex networks have gained significant attention (Dai et al., 2019; Wang et al., 2020; Huang et al., 2021; Yin et al., 2022; Chen et al., 2022b; Fügenschuh and Fu, 2023; Xu et al., 2024). The majority of these studies focus on understanding how the dynamics of information and disease across different layers influence disease spread. However, limited research has explored intervention strategies in multi-layer settings (Fügenschuh and Fu, 2023), highlighting the need for further investigation into control performance

within the complex interplay between information and disease spreading across different network structures.

Considering the different channels of information diffusion and recognizing that social media platforms are key venues for opinion exchange with a growing influence on public opinion, we aim here to evaluate the effectiveness of control strategies in a multi-layer network model [1]. In particular, we consider a multiplex network consisting of distinct layers for information diffusion and disease spread. We aim to examine the performance of the proposed intervention campaigns in a two-layer network, addressing several key facets. Given that information and opinions can spread across different types of social networks , e.g., virtual and face-to-face interactions, unlike disease transmission which typically requires physical contact, a key question arises: which layer should be targeted for intervention? Thus, we explore whether an intervention based on information from the virtual layer is more effective than one based solely on the physical layer, which overlooks virtual contacts. We also investigate how the correlation between the two layers influences the resulting dynamics. In addition, we investigate the impact of varying topological configurations of these layers.

Specifically, this chapter makes the following contributions to address research questions 4a and 4b detailed in Section 1.1. We examine two intervention schemes: intervention targeting virtual networks and intervention targeting physical networks to determine which type of social network is more effective in curbing the negative diffusion and, consequently, disease transmission. We conduct a factorial analysis exploring several configurations that consider the topology of each layer, the overlap between layers, and the choice of target layer.

The chapter is structured as follows. Section 7.2 describes the model and the construction of the multi-layer network. Section 7.3 presents the results. Specifically, Section 7.3.1 shows the results of intervention through virtual networks, Section 7.3.2 presents the results of intervention through physical networks, and Section 7.3.3 provides a comprehensive comparison and discussion of both approaches. Finally, Section 7.4 summarizes the main findings.

7.2 Model Description and Methods

We consider a multiplex network with two layers, in which one layer represents information (opinions) diffusion and the other represents disease spread. The information layer represents the virtual network (e.g. X or Facebook) and is denoted by G_v. The disease spread layer represents face-to-face interactions and is denoted by

[1] Throughout this work, we use the terms 'multi-layer' and 'multiplex' networks interchangeably.

G_d, referred to as the physical layer. A schematic representation of a multiplex network with two layers is shown in Figure 7.1.

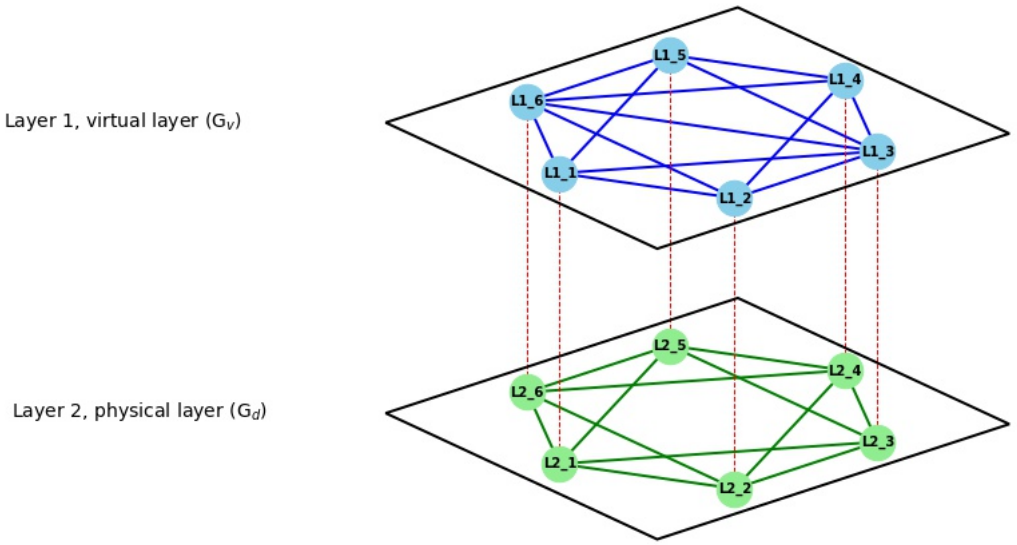

FIGURE 7.1: Schematic representation of a multiplex network consisting of two layers, a virtual layer (G_v) and a disease propagation layer (G_d), with each node linked to its counterpart across layers. Note that the virtual layer contains more connections than the physical layer.

Despite the extensive research on multiplex network models, there is no standardized framework for constructing layer configurations in the existing literature. For example, Dai et al. (2019) used Erdos-Rényi (ER) networks for all layers, Wang et al. (2020) considered a combination of square lattice and ER networks, Yin et al. (2022) used a configuration of Barabási– Albert and ER networks, and Fügenschuh and Fu (2023) considered Barabási– Albert and square-lattice networks.

Here, as many real online social networks exhibit a power-law degree distribution (Mislove et al., 2007; Sadikov and Martinez, 2009), we consider a scale-free network for the virtual layer. For the physical layer, we retained the small-world network previously used in this thesis. Furthermore, since social media often mirrors physical contact networks while also introducing virtual friendships, the virtual layer is constructed by extending the physical layer with additional edges. This approach has also been adopted in the existing literature (Huang et al., 2021). Specifically, the physical network is replicated and transformed into a scale-free structure by adding new edges, K_v, using the Barabási–Albert model to build the virtual network. Empirical support for these modeling assumptions can be found in (Mislove et al., 2007; Said et al., 2019), which highlight the presence of both small-world and scale-free properties in social media networks.

Furthermore, two intervention schemes are considered virtual-based and physical-based interventions. In virtual-based interventions, campaigns target

individuals in the virtual layer, leveraging connection information from that layer. In practice, this corresponds to applying intervention on social media platforms, utilizing social contact data derived from these platforms. In contrast, physical-based interventions target individuals in the physical layer, relying on connection information from the physical layer, while information propagation occurs in the virtual layer. This represents face-to-face contact networks such as schools, workplaces, and others.

7.3 Results

This section presents results on epidemic size obtained from applying the intervention campaigns proposed in previous chapters to the multi-layer network. We begin by evaluating the campaigns applied to the virtual layer (Section 7.3.1). Next, we assess the control achieved when the intervention is applied to the physical layer (Section 7.3.2). Finally, we present a comprehensive comparison and analysis of these intervention schemes (Section 7.3.3).

Since epidemic size decreases monotonically as the budget increases, we focus solely on the equal-budget scenario in this section. For the physical layer, we maintain the same setting used in previous chapters, considering a small-world network of size N = 5000, with an average degree $\langle k \rangle$ = 10 and a rewiring probability p = 0.01. The virtual network layer is a scale-free network, with other configuration settings presented in the following sections. We preserve the same diffusion parameters used in previous chapters. The opinion formation threshold is set to $\theta = 2$, while the SIR model parameters are set as follows: an infection rate $\beta = 0.1$, a recovery rate $\gamma = 0.1$, and a seed set $I_0 = 1$. Unless otherwise stated, results are the average of 50 realizations. In all figures presented, the error bars correspond to the 95% confidence intervals.

7.3.1 Control Through the Virtual Layer

To understand how information diffusion in the virtual layer influences disease spread in the physical layer, and to explore the extent to which controlling the virtual layer can curb disease spreading, we present here the final epidemic size when control is applied to the virtual layer. The results are presented for different intervention timings t_s^+, considering early intervention at $t_s^+ = 100$ with 1% of the population has adopted an anti-vaccine opinion, and late intervention at $t_s^+ = 145$ with 4% of the population has adopted an anti-vaccine opinion. These intervention timings are consistent with those used in Chapter 6 for scale-free networks, as opinion diffusion here occurs on a scale-free network.

Figure 7.2 shows the final epidemic size obtained across all intervention campaigns target individuals in the virtual layer G_v. As can be seen in the figure, with early intervention at $t_s^+ = 100$ (Fig. 7.2 (A)), all campaigns perform reasonably well in this early intervention, except for the DynRandT campaign. In addition, the behavior of the intervention campaigns follows a similar control pattern to that observed in scale-free networks (see scale-free network in Fig. 6.2), despite the disease spreading within small-world structures. More specifically, the best performance is achieved using the Min_NK$_{wt}$ campaign (red color) as well as DynAdvLocT (blue color) [2].

With delayed intervention at $t_s^+ = 145$ and a higher proportion of anti-vaccine opinion adopters (Fig. 7.2 (B)), a similar trend is observed; however, the effectiveness of control is reduced in this scenario, resulting in a larger epidemic size. In this scenario, the Min_NK campaign (black bar) shows reduced effectiveness, while the the Min_NK$_{wt}$ campaign performs significantly better. This contrasts with their performance in a purely scale-free network, where both demonstrate similar performance (see Fig. 6.2 for SF networks). However, in the small-world network, the Min_NK$_{wt}$ campaign outperforms the Min_NK campaign (see Fig. 6.2 for SW networks). The improved performance here is therefore attributed to the inherent small-world nature of the virtual network. The DynAdvLocT campaign shows marginally better results than the optimal control (Min_NK$_{wt}$) campaign.

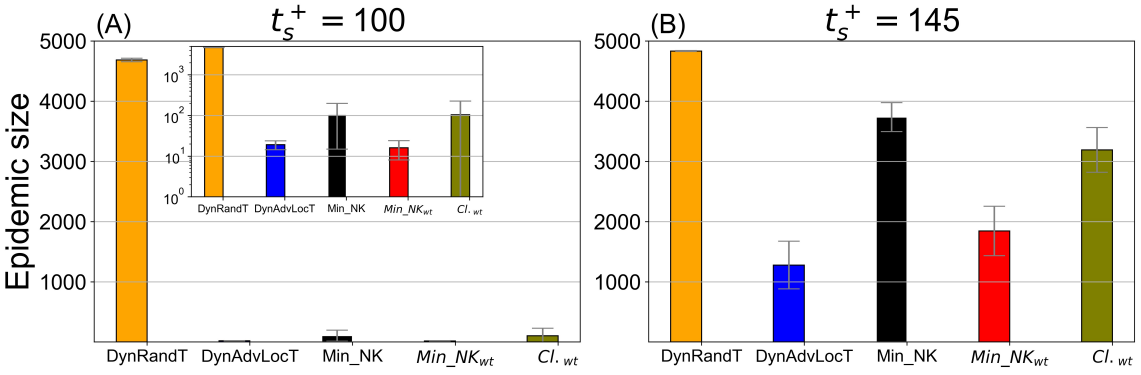

FIGURE 7.2: Final epidemic size for intervention campaigns targeting the virtual layer G_v: (A) early intervention at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 38 $\pm$3, and (B) late intervention at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 175 $\pm$16. The average degree $\langle k \rangle$ for $G_v = 14$ and $G_d = 10$. Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$. For DynAdvLocT campaign $\zeta = 10$ and $Z = 100$.

As for the cluster-based campaign (Cl.$_{wt}$), it is no longer efficient with delayed intervention at $t_s^+ = 145$ despite the network's inherent high clustering (see olive bars in Fig. 7.2 (B)), consistent with what was observed in the scale-free network (Fig. 6.2). As previously demonstrated, the clustering pattern of anti-vaccine opinion adopters in a scale-free network makes the cluster-based approach less effective in this

[2] Additional details on the dependency of the DynAdvLocT campaign on its parameters Z and ζ can be found in Appendix C.1.

topology. This is due to the formation of large, often single, anti-vaccine communities from early stages, which makes it difficult for the campaign to contain their persistent growth. This raises the question: Can controlling clusters in the physical layer provide more effective control than targeting virtual communities? In the next section, we examine the control efficiency across all campaigns when targeting the physical layer.

7.3.2 Control Through the Physical Layer

Here, we assess the performance of intervention campaigns targeting individuals in the physical layer G_d. In this scenario, virtual contacts per individual are unavailable, as we only have information about physical contacts. This means we are evaluating the campaigns with missing contact data. Therefore, this scenario will enable us to examine whether the lack of virtual contact information impacts the control and mitigation of information diffusion, and whether the spread of information becomes difficult to control through the physical layer.

Figure 7.3 presents the final epidemic size obtained for early intervention at $t_s^+ = 100$ and late intervention at $t_s^+ = 145$. As shown in the figure, with early intervention at $t_s^+ = 100$ (Fig. 7.3 (A)), the campaigns can effectively controls information diffusion despite the missing contact data from the virtual layer, thereby reducing the epidemic size. The DynAdvLocT and $Cl._{wt}$ campaigns demonstrate the best performance, while the Min_NK$_{wt}$ campaign achieving comparable results. However, even with the most effective campaign in this scenario, approximately 20% of the population becomes infected, in contrast to interventions in the virtual layer, where the majority of the population was protected at the same intervention timing, see Fig. 7.2 (A).

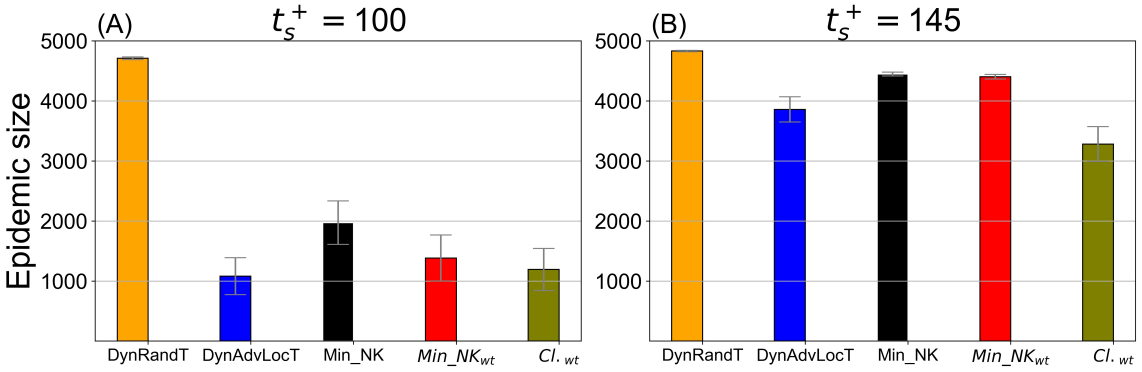

FIGURE 7.3: Final epidemic size for intervention campaigns targeting the physical layer G_d: (A) early intervention at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters 38 ± 3, and (B) late intervention at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters 175 ± 16. The average degree $\langle k \rangle$ for $G_v = 14$ and $G_d = 10$. Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

With late intervention at $t_s^+ = 145$ (Fig. 7.3 (B)), the performance of all campaigns deteriorates significantly, leading to the infection of the majority of the population. In this scenario, the best control is achieved by the cluster-based approach $Cl._{wt}$; however, even with this strategy, more than half of the population becomes infected. In contrast, intervening at this stage through the virtual layer results in significantly greater mitigation of negative influence and containment of disease spread (Fig. 7.2 (B)).

As observed, the absence of contact data from the virtual layer negatively impacts the effectiveness of intervention campaigns. While there is potential to control information diffusion and achieve better outcomes through the physical layer with early intervention, interventions through the virtual layer demonstrate superior control. This raises the question: To what extent does the control effectiveness differ between interventions applied to the virtual and physical layers? In the next section, we quantify and compare the control effectiveness achieved through interventions in both layers.

7.3.3 Control: Virtual vs. Physical Layers—Which Matters Most?

Here, we compare and quantify the variation in control achieved through interventions across virtual and physical layers. To illustrate the extent of this variation, Fig. 7.4 presents a comparison of the resulting epidemic size from virtual-based ($G_{\bar{v}}$) and physical-based (G_d) interventions in the scenario discussed above. Figure 7.4 (A) shows results for early intervention at $t_s^+ = 100$, while Fig. 7.4 (B) shows results for late intervention at $t_s^+ = 145$ across all proposed campaigns. As shown in the figure, there is a significant difference in control performance between the two intervention approaches across all intervention timings. A substantial level of control can be achieved through virtual-based intervention, while a notable decline in control is observed with physical-based intervention. The overlap between virtual and physical networks enables virtual-based interventions to incorporate social contacts from different types of relationships, including both virtual and physical contacts. In contrast, physical-based interventions are limited to face-to-face interactions only. With early intervention at $t_s^+ = 100$, targeting based on the virtual layer achieves approximately 98% greater reduction through the DynAdvLocT campaign compared to targeting based on the physical layer, and around 99% improvement for the Min_NK$_{wt}$ campaign. Notably, similar performance enhancement is observed even with delayed intervention, yielding a 67% improvement for DynAdvLocT and 58% for Min_NK$_{wt}$.

Despite the physical network consisting of a large number of connections per agent, we observed that the absence of a few virtual-layer contacts through physical-based intervention significantly degrades performance. Note that, in the current scenario,

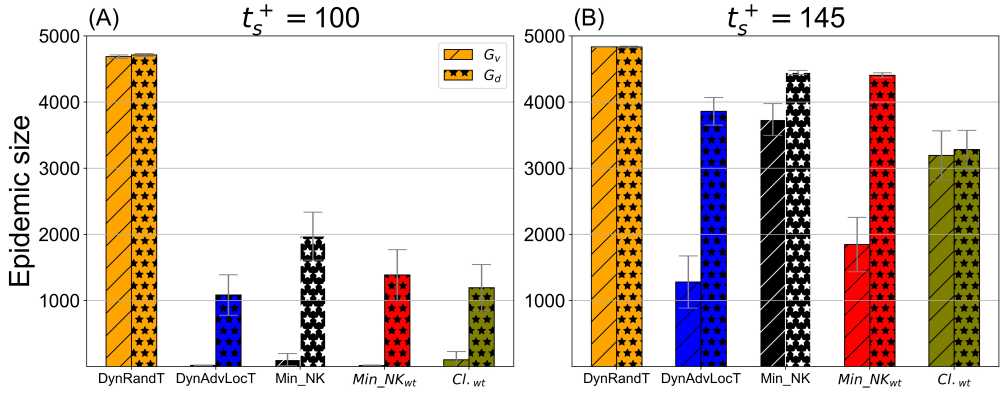

FIGURE 7.4: Comparison of the epidemic size resulting from virtual-based (G_v) and physical-based (G_d) intervention campaigns. The average degree $\langle k \rangle$ for $G_v = 14$ and $G_d = 10$. G_v is a scale-free network extended from the physical layer. Intervention campaigns starting at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 38 $\pm$3 (left panel), and starting at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 175 $\pm$ 16 (right panel). Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

the average degree of the virtual layer is $\langle K_{G_v} \rangle = 14$, and that of the physical layer is $\langle K_{G_d} \rangle = 10$. The average degree represents the average number of contacts per individual in a network (see Chapter 2, Section 2.1), with higher values indicating more socially connected populations and faster information propagation. This discrepancy results in a marked reduction in control effectiveness through physical-based intervention the longer the intervention is delayed. To quantify this impact, we next reduce the degree in the virtual layer to narrow the gap in the average degree between the two layers and examine whether this also results in a substantial effect on performance.

Figure 7.5 shows a comparison of the epidemic size for virtual-based and physical-based interventions, with the average degree of the virtual layer is $\langle K_{G_v} \rangle = 12$, and that of the physical layer is $\langle K_{G_d} \rangle = 10$. Figure 7.5 (A) shows the results for early intervention at $t_s^+ = 100$, while Fig. 7.5 (B) shows the results for late intervention at $t_s^+ = 145$ across all intervention campaigns. Similarly, as shown in the figure, the virtual-based intervention outperforms the physical-based intervention with a noticeable difference for all campaigns, despite the latter having fewer missing virtual contacts than in the previous scenario. This performance difference remains approximately the same in both the early and late intervention scenarios, with the virtual-layer approach delivering substantially better performance through DynAdvLocT and Min_NK$_{wt}$ campaigns. These campaigns achieve reductions of 94% and 98% for the DynAdvLocT and Min_NK$_{wt}$ campaigns, respectively, at $t_s^+ = 100$, and 90% and 69% at $t_s^+ = 145$. Moreover, the overall gain in control is greater, resulting in a smaller epidemic size compared to the previous scenario in Fig. 7.4, due to the lower density of the virtual network. For example, the resulting epidemic size

through DynAdvLocT campaign (blue bars) at the intervention time $t_s^+ = 145$ is noticeably smaller when comparing Figs. 7.4 (B) and 7.5 (B).

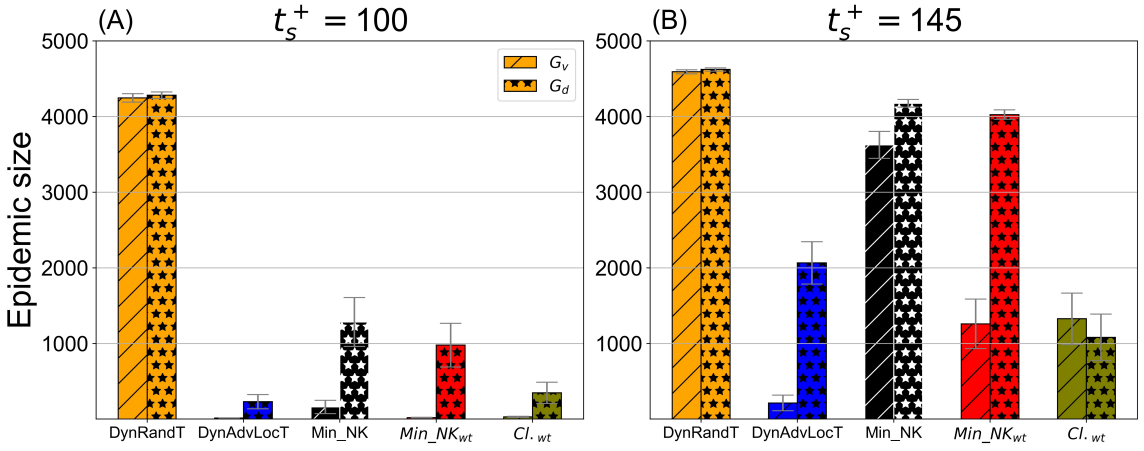

FIGURE 7.5: Comparison of the epidemic size resulting from virtual-based (G_v) and physical-based (G_d) intervention campaigns. G_v is a scale-free network extended from the physical layer. The average degree $\langle k \rangle$ for $G_v = 12$ and $G_d = 10$. Intervention campaigns starting at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 35 ± 3 (left panel), and starting at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 126 ± 9 (right panel). Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

This underscores the importance of considering the contact information from the virtual layer, where information and opinions are exchanged more extensively and widely than in the physical layer, to achieve high control efficiency. Furthermore, the importance lies not only in the extent of missed virtual contacts but also in their connectivity, highlighting that the topology of the virtual layer is a crucial factor to consider. Note that, in previous scenarios, the virtual layer exhibits a scale-free network, where information spreads more rapidly and extensively. These findings suggest that even a small difference between the virtual and physical layers can substantially reduce control, not only due to missing virtual data but also due to virtual structural connections.

To further explore the role of virtual layer topology, we consider a scenario where both layers exhibit the same type of structure. Figure 7.6 presents the results of a configuration where both layers exhibit a small-world network, with the virtual layer extending from the physical layer. Although this assumption might be unrealistic given that existing literature demonstrates many online social networks typically exhibit a scale-free structure (as discussed earlier), we adopt this assumption here to strengthen our understanding and observations. In this scenario, the average degree of the virtual layer is $\langle K_{G_v} \rangle = 12$, and that of the physical layer is $\langle K_{G_d} \rangle = 10$. As can be seen in Fig. 7.6, both intervention schemes, virtual-based and physical-based, provides similar performance. This reinforces our observation on the significance of the information layer structure. Since both layers share the same structure and the

majority of contacts are included in the physical-based intervention, the absence of virtual contact data does not impact control performance. As a result, negative influence diffusion is effectively controlled and contained with our campaigns.

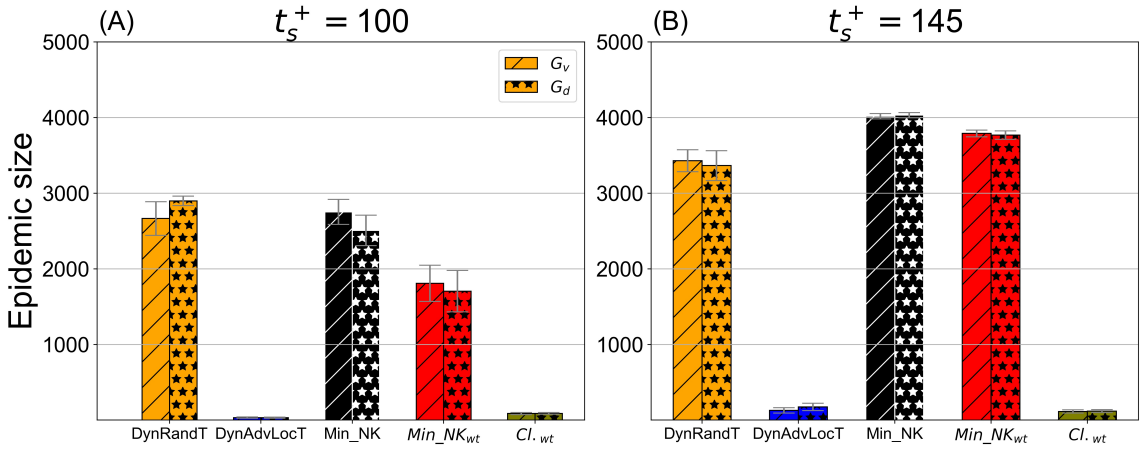

FIGURE 7.6: Comparison of the epidemic size resulting from virtual-based (G_v) and physical-based (G_d) intervention campaigns. G_v is a small-world network extended from the physical layer. The average degree $\langle k \rangle$ for $G_v = 12$ and $G_d = 10$. Intervention campaigns starting at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 37 ± 2 (left panel), and starting at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 131 ± 8 (right panel). Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

Next, what if there is no correlation between the two network layers? Does the virtual layer still matter the most? Here, we consider a scenario where no correlation exists between the two layers, relaxing the assumption of extended layers and generating distinct networks for each layer. In this configuration, we maintain the same setting considered before for the average degree and network structures for both layers. The virtual layer is modeled as scale-free network, extended from a small-world network that differs from the physical layer, with an average degree $\langle K_{G_v} \rangle = 12$, while the physical layer is modeled as small-world network with average degree $\langle K_{G_v} \rangle = 10$. To assess the impact of each layer, we again implement virtual-based and physical-based interventions at two intervention timings. The goal is to determine which layer has a greater influence on the outcomes of the intervention in this uncorrelated multilayer network scenario and to identify whether the virtual or physical layer plays a more significant role in controlling information diffusion[3].

Figure 7.7 shows the results for this uncorrelated layers scenario. We observe that early intervention results in efficient control for both virtual-based and physical-based interventions, with the exception of the dynamic random campaign (Fig. 7.7 (A)). In the case of late interventions (Fig. 7.7 (B)), the distinction in control efficiency is more pronounced, with virtual-based intervention outperforming control through the

[3]we also evaluate a scenario where the virtual layer is a scale-free network, not extended from a small-world network, and the results align with the findings presented here (see Fig.C.2 in the Appendix).

physical layer, except for the Min_NK and cluster-based ($Cl._{wt}$) campaigns. Controlling physical communities is much more efficient than controlling virtual communities, as expected, which is also consistent with the findings in Fig. 7.5 (B) with marginal improvements. For the Min_NK campaign, minimizing the number of anti-vaccine individuals is more efficient in the physical layer than in the virtual layer, although it is not the most effective campaign. In contrast, the Min_NK$_{wt}$ campaign, which focuses on minimizing negative influence while fostering positive influence, performs better in virtual-based targeting. Similar performance also observed for DynAdvLocT campaign. These results further support our observation that the virtual layer is the most important, and controlling the virtual layer is more crucial than intervening in the physical layer.

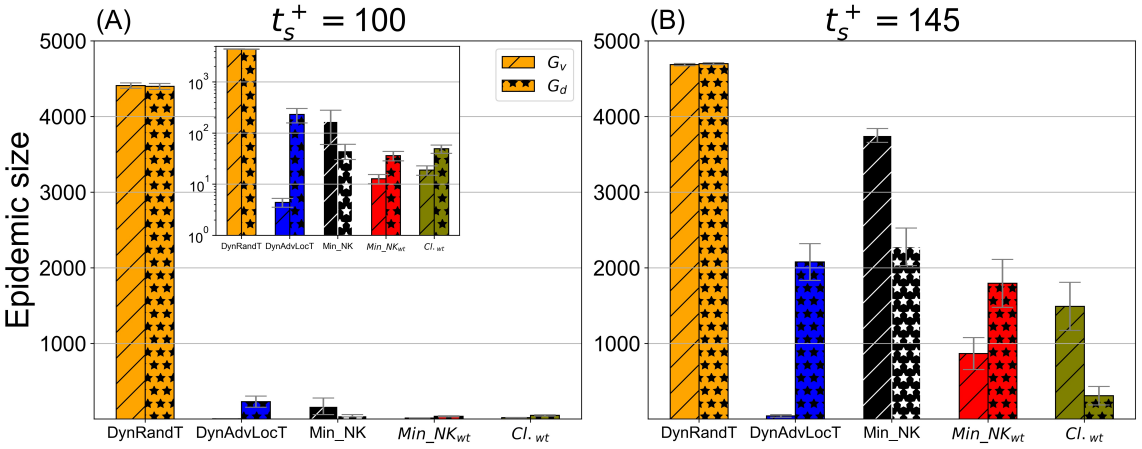

FIGURE 7.7: Comparison of the epidemic size resulting from virtual-based (G_v) and physical-based (G_d) intervention campaigns. Uncorrelated layers, (G_v) is a scale-free network extended from a small-world network. The average degree $\langle k \rangle$ for $G_v = 12$ and $G_d = 10$. Intervention campaigns starting at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 39 ± 3 (left panel), and starting at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 133 ± 10 (right panel). Parameters values: $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

To provide an overview of intervention performance across different layers, Table 7.1 summarizes the relative change in epidemic size when using virtual-based versus physical-based interventions for each campaign across both correlated and uncorrelated layer settings, considering early and late interventions. In the table, positive values indicate that virtual-based interventions result in smaller epidemic sizes than physical-based interventions, and vice versa for negative values. As shown in the table, virtual-based interventions generally outperform physical-based interventions across most campaigns, with their effectiveness improving when applied earlier. The notable exception to this trend is the cluster-based campaign Cl.$_{wt}$, where physical-based interventions outperform virtual-based ones in both correlated and uncorrelated settings. These results highlight the critical role of timing,

TABLE 7.1: Relative change in epidemic size between virtual-based and physical-based interventions in correlated layers (Figure 7.5) and uncorrelated layers (Figure 7.7). Negative values indicate that the epidemic size achieved with a physical-based intervention is smaller.

Campaign	Correlated (Early)	Correlated (late)	Uncorrelated (Early)	Uncorrelated (late)
DynRandT	0.8 %	0.7 %	-0.3 %	0.3 %
DynAdvLocTT5	94 %	90 %	98 %	98 %
Min_NK	88 %	13 %	-274 %	-65 %
Min_NK$_{wt}$	98 %	69 %	65 %	52 %
$Cl._{wt}$	91 %	-23 %	62 %	-382%

intervention type (virtual vs. physical), and the overlap between layers of system dynamics in achieving optimal epidemic control.

7.4 Summary

In this chapter, we increase the realism of our model by considering different opinion exchange platforms through multiplex networks, where we assume separate layers for information diffusion and disease spread. We conducted a factorial analysis considering three key factors: the configuration of the layers, the correlation between layers, and which layer to target. Through extensive experiments, several key insights have been identified. First, the behavior of campaigns in multi-layer networks remains consistent with their behavior in single-layer networks. The choice of intervention strategy should align with the structural characteristics of the virtual layer, where information spreads widely. For example, strategies effective in single-layer scale-free networks maintain similar efficacy in multi-layer settings where information propagates in a scale-free structure, despite disease spread occurring in a different structure.

In physical-based interventions that consider only face-to-face interactions, performance can be affected and degrade due to the absence of virtual data connections for agents. The extent of this impact depends on the structure of the virtual network. Our findings reveal that in networks where the virtual layer follows a scale-free structure, missing virtual contact data significantly diminishes performance, particularly during the widespread propagation of negative influence in late interventions.

Further, the findings demonstrate that virtual-based interventions outperform physical-based interventions across most campaigns. The notable exception to this trend is observed in the cluster-based campaign (Cl.$_{wt}$), particularly with late interventions, where controlling physical anti-vaccine communities perform better

than controlling virtual anti-vaccine communities, regardless of whether there is a correlation between the network layers.

Chapter 8

Conclusions and Future Work

This thesis addresses the problem of mitigating the propagation of vaccine misinformation through strategic interventions, with the primary goal of curbing the spread of vaccine-preventable diseases. It explores how external control targeting vaccine-related opinion dynamics can influence the trajectory of epidemics. By addressing several key challenges identified in the field, this study advances our understanding and provides new perspectives and practical insights. Below, we summarize the main contributions and findings of this research and highlight its limitations as well as potential directions for future work.

8.1 Conclusion

In this thesis, we propose several methods to disseminate pro-vaccine information and encourage vaccine uptake to counter the spread of anti-vaccine attitudes, with the ultimate goal of suppressing disease spread. Through extensive numerical experiments, we demonstrate the effectiveness of these campaigns in impeding the flow of anti-vaccine attitudes and, consequently, reducing the spread of epidemics. However, different methods of distributing positive influence have varying impact on anti-vaccine opinion dynamics and the structure of anti-vaccine communities, which in turn affect the extent of epidemic suppression. We find that several factors play crucial roles in the effectiveness of these strategies, including the allocated positive budget, the timing of interventions, the underlying network structure, and the extent of the spread of negative opinions prior to the interventions.

In Chapter 3, we propose several heuristics based on two paradigms: network topology and vaccine-related local information, and two control schemes: static and dynamic. We design adaptive strategies based on local information about vaccine attitudes, exploring the trad-off between blocking negative influence diffusion and

fostering positive influence diffusion. We demonstrate that, first, targeted campaigns that select a subset of the population based on certain criteria have been found to be more effective in suppressing the epidemic compared to the random campaign, particularly in scenarios with high social influence rate. Furthermore, the centrality-based strategy, which targets central nodes, and local information-based strategies, which target individuals susceptible to negative influence, exhibit superior performance. Second, the key feature of static campaigns, with the fixed positions of the targets, is their ability to create barriers within the network that effectively disrupt the clustering of the anti-vaccine individuals. In contrast, dynamic campaigns can reach and influence a broader segment of the population by continually targeting individuals who are vulnerable to negative influence from their social contacts. Third, the dynamic control strategies, due to their adaptive nature, have been found to be more effective in suppressing the epidemic than the static control approach. Among the dynamic campaigns, the DynAdvLocT strategy shows the best performance, as it not only protects the most vulnerable agents from negative influence but also amplifies the spread of positive influence. Finally, other factors also play a significant role in the performance of each campaign, including the target set size, the updating time for dynamic campaigns, and the social influence strength. These findings solved research questions 1a,1b, and 1c detailed in Section 1.1.

In Chapter 4, we advance our investigation into efficient methods to counter the diffusion of anti-vaccine influence by leveraging optimization techniques. In particular, we frame the problem as an optimization problem, exploring two distinct optimization objectives: minimizing the number of anti-vaccine opinion adopters and maximizing the number of pro-vaccine opinion adopters. We find that, first, minimizing the number of anti-vaccine opinion adopters (Min_NK) is more effective in containing the diffusion of negative influence than maximizing pro-vaccine opinion adopters. This is because the minimization strategy targets agents in close proximity to negative areas, enabling the seeding of positive influence in high-risk regions of the network. Consequently, it helps disrupt connectivity and limit the expansion of emerging anti-vaccine communities, resulting in a significant reduction in epidemic size. In contrast, the maximization strategy operates primarily near existing pro-vaccine regions, often at a distance from negative regions. As a result, it fails to curb the spread of anti-vaccine influence and the growth of anti-vaccine communities, resulting in poor containment of the epidemic. Second, the effectiveness of the minimization campaign diminishes the longer the intervention is delayed. This is due to the early formation of anti-vaccine clusters, which continue to grow despite campaign efforts to contain them, particularly under limited budget scenarios. These findings solved research questions 1d and 2a detailed in Section 1.1.

In Chapter 5, we investigate the problem of containing the growth and expansion of existing anti-vaccine communities. We demonstrated in the previous chapters that the

performance of the best-performing campaigns proposed, i.e., the DynAdvLocT and the Min_NK, decrease the longer the interventions are delayed. Therefore, in this chapter, we extend our optimization framework and develop two cluster-based strategies. These strategies identify and target critical regions in the network by targeting agents bridging multiple existing clusters. We find that, first, these are efficient containment methods as they mitigate the emergence of large-scale clusters by shielding critical regions, resulting in a significant reduction in epidemic size. Second, prioritizing the control of larger anti-vaccine clusters over smaller ones yields even better results. These findings solved research question 2b detailed in Section 1.1.

In Chapter 6, we study the influence of network structure on the performance of the proposed strategies, identifying the most effective approaches for each network type, the limitations of each strategy, and the underlying reasons for these outcomes. In this regard, taking into account the heterogeneity of different networks, we further optimize the minimization campaign by integrating knowledge about the contact structure of agents. Under similar conditions, i.e., the same available budget, number of anti-vaccine individuals prior to intervention, average degree, and network size, we observe significant variations in control performance. First, the results reveal that network structure significantly affects both mitigation efforts and the level of control achieved, and there is no single campaign that works effectively for all topologies, particularly for the case of late interventions. Second, we find that minimization campaigns, particularly the optimized campaign (Min_NK$_{wt}$), are effective with early interventions across all network structures. However, their effectiveness diminishes the longer the intervention is delayed. Third, in late intervention scenarios and with a greater dominance of negative influence, negative spread can be effectively controlled by shielding critical regions with the cluster-based campaigns in small-world and spatial networks, which significantly reduces the resulting epidemic size. However, this approach is less effective in scale-free and random networks. This behavior stems from the initial structures of existing anti-vaccine communities and the evolving dynamic nature of these clusters in scale-free and random networks, as they tend to produce large clusters from the early stages of their emergence, making their containment increasingly challenging with limited budget. These findings solved research questions 3a and 3b detailed in Section 1.1.

In Chapter 7, we address the robustness of control performance in a more realistic scenario, where information diffusion and disease transmission occur on different social networks, capturing the diversity of communication patterns through a multi-layer network. We conduct a factorial analysis considering three key factors: structural differences between layers, their correlation, and the targeted layer. We find that, first, the behavior of campaigns in multi-layer networks is consistent with that in single-layer networks, and the choice of intervention strategy should align with the structural characteristics of the virtual layer where information spreads most densely.

Second, incomplete social contact data can undermine the effectiveness of control campaigns. This is particularly relevant in scenarios when interventions target the physical network, i.e., face-to-face connections, where contact data from the virtual network is missing. However, the extent of this degradation depends on the structure of the virtual network where it is particularly significant when this layer has a scale-free structure, most prominently during the widespread propagation of negative influence in late interventions. Finally, virtual-based interventions generally outperform physical-based interventions across most campaigns, except for the cluster-based campaign (Cl.$_{wt}$), where controlling physical anti-vaccine communities performs better than controlling virtual anti-vaccine communities, regardless of layer correlation.

These findings advance our theoretical understanding of effective approaches to deploy available resources to efficiently control the coupled dynamics of information and disease spread. Despite the real world being much more complex than our assumptions, key insights have been drawn systematically that can inform practical interventions. Our findings suggest that influencing anti-vaccine opinion dynamics has a significant positive impact on controlling the spread of epidemics. In addition, effectively controlling disease transmission requires not only mitigating the prevalence of anti-vaccine opinions but also disrupting the growth of their clusters.

We explore different control approaches and find that responding to the evolving dynamics of negative diffusion in the network provides effective containment. Specifically, focusing on regions in proximity to anti-vaccine attitudes to block its diffusion while simultaneously fostering the spread of pro-vaccine influence is an effective approach. We also demonstrate the importance of strategically promoting vaccination in critical network regions that exist between anti-vaccine clusters to initiate protective barriers that block the expansion of negative areas, particularly in scenarios where initial opportunities for early action have been missed. However, while identifying these regions is easy in some network structures, it might be challenging in others. Although perfect practical implementation may be challenging, these insights could inform actionable heuristics for public health campaigns. For example, focusing interventions on individuals or areas positioned between clusters of negative opinions could help maximize the impact of campaigns.

Translating these theoretical insights into practice is feasible, requiring parameter calibration using empirical observations including data on the network structure and vaccine attitudes. However, this involves several practical challenges. Our model is designed to be flexible and adaptable to real-world needs, although it can be improved further. In principle, the model's parameters can be calibrated with real-world data to represent localized dynamics, such as the spread rates of positive and negative public health information, the distribution of anti-vaccine, neutral, and pro-vaccine individuals, and the structure of the network. In addition, our model

operates on the assumption of comprehensive knowledge regarding vaccine-related attitudes within the social network. Full knowledge of the attitude of every individual is not realistic in practice; however, obtaining some information is feasible from social media platforms. For instance, empirical studies leveraging data from online social networks have applied sentiment analysis and machine learning algorithms to categorize individuals' attitudes into distinct states such as pro-vaccine, anti-vaccine, and neutral toward vaccination (Salathé and Khandelwal, 2011; D'Andrea et al., 2019; Abd Rahim and Rafie, 2020; Yousefinaghani et al., 2021). Additionally, public attitudes toward vaccines have been explored through the analysis of real-time, spatial-temporal, and socio-demographic data from social media, unveiling the spatial distribution of these attitudes (Hu et al., 2021; Umair and Masciari, 2023; Cheng et al., 2023). Such platforms and methods enable real-time, socio-geographic monitoring of public attitudes.

While vaccine attitudes can be obtained from such platforms, acquiring accurate data might be challenging and error-prone, and might not fully capture the complexity of the real world. We have addressed this practical challenge by examining the impact of data noise on campaigns' performance and demonstrated that, although noise can influence the effectiveness of intervention campaigns, several strategies exhibit strong resilience across various scenarios, particularly when sufficient budget is available. Notably, for optimization-based strategies, the results suggest that the presence of noise in the data might actually be beneficial in some cases, as it can reduce the system's sensitivity by broadening the targeting pool for these campaigns.

Another key practical challenge is determining which social venues should be targeted to effectively manage anti-vaccine diffusion and epidemic outbreaks, given the diversity of social network platforms, including social media and face-to-face interactions. Virtual networks enable rapid dissemination of information and opinions, reaching a broader segment of society that can influence vaccine uptake and epidemic dynamics in the physical space. This is reflected in prior empirical studies that have demonstrated a correlation between vaccine attitudes and vaccination rates (Pierri et al., 2022; Cheng et al., 2023). Through a multi-layer network model, we demonstrate that interventions targeting virtual networks to control opinion propagation can effectively mitigate epidemic outbreaks in the physical network. These interventions also exhibit superior efficiency compared to interventions targeting physical networks in most scenarios.

8.2 Limitations and Future Work

Although the model provides valuable insights, several inherent limitations should be acknowledged to fully understand its scope and applicability. First, the presented

results are limited to the assumption of committed agents who do not change their opinions once adopted. However, the proposed heuristics could be extended to a broader range of opinion diffusion models and integrated into models that allow opinions to switch back and forth, such as the voter model and epidemic models like the SIS model. In addition, empirical studies have highlighted that factors such as education, gender, and other socioeconomic characteristics can influence individual decisions (Morales et al., 2022). Therefore, the model could be further extended to incorporate sociodemographic data to capture how these sociodemographic characteristics shape agent preferences and beliefs.

Furthermore, opinions in our model are not influenced by disease spread, as we address the diffusion processes of information and disease separately (in line with similar studies, e.g., (Salathé and Bonhoeffer, 2008; Campbell and Salathé, 2013)). Information spread is assumed to occur before the onset of disease, aligning with patterns typically observed in vaccine-preventable childhood diseases. The proposed campaigns can be implemented and evaluated in scenarios where information dissemination and disease transmission can occur concurrently, such as the spread of information related to newly developed vaccines during ongoing outbreaks like Covid-19. Similarly, the model can be extended to capture scenarios in which opinion adoption could also be influenced by disease dynamics such as awareness of infected neighbors.

In this study, we focus on measles-like diseases, where vaccination provides full immunity and vaccinated individuals neither acquire nor transmit the infection. Future work could extend this approach and explore the applicability of the proposed strategies to diseases with different characteristics, such as Covid-19 or flu-like illnesses, taking into account factors such as vaccine efficacy, partial immunity, and the need for repeated vaccination. Similarly, pathogens exhibit varying properties such as differences in transmission rates (across pathogens or even mutations), infection severity, mortality, and others. Further extension could incorporate pathogen-related factors and explore the effectiveness of interventions on disease dynamics.

In addition, the proposed strategies, though efficient, can be further improved and extended. For instance, incorporating additional behavioral or epidemiological factors could enhance their effectiveness. Optimization-based campaigns could also be validated using real-world social media, where substantial noise is expected, enabling a further discussion of campaigns performance and resilience, as well as identifying opportunities for further enhancement. Furthermore, in the virtual layer of the multi-layer model, we assume that individuals are influenced only by content from their directly connected friends. However, most social media platforms employ engagement-based algorithms (Narayanan, 2023) that can show content tailored to users' interests, including content from users outside their network. Future extensions of this model could explicitly incorporate the effects of such algorithms.

Appendix A

Appendix to Chapter 3

A.1 Comparison of Different Centrality Measures

In this section, we present the epidemic sizes obtained using different centrality metrics applied in our campaign to select the target set: betweenness centrality and degree centrality. As illustrated in Fig A.1 below, the differences between the two metrics are minimal, with betweenness centrality showing slightly better performance.

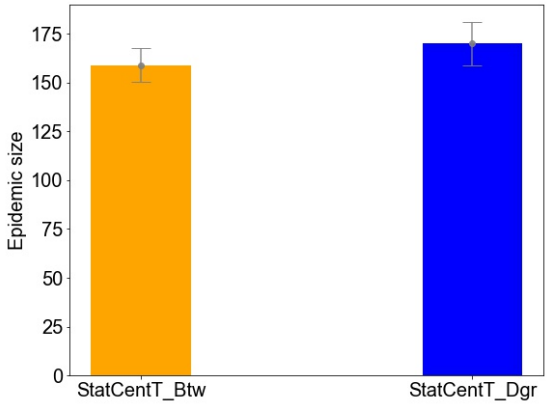

FIGURE A.1: Comparison of epidemic size using different centrality measures: betweenness centrality (StatCentT_Btw) and degree centrality (StatCentT_Dgr). For all campaigns, $\mu^- = \mu^+ = 0.001$, the social rate is $\omega = 0.006$, the target set size is $T = 500$, and $\tau = \infty$. The results are the average of 100 simulations.

A.2 Disease Evolution

Figure A.2 below illustrates the evolution of the disease after implementing the positive campaign in various ways. It shows the evolution of both the number of

infected cases and the number of recovered cases. As shown, the peak of infected cases and the duration of the disease are significantly reduced when using dynamic campaigns, particularly DynRandT and DynAdvLocT, compared to the baseline StatRandAll, where the entire population is targeted randomly. These dynamic campaigns also outperform static campaigns, although the latter still perform better than StatRandAll.

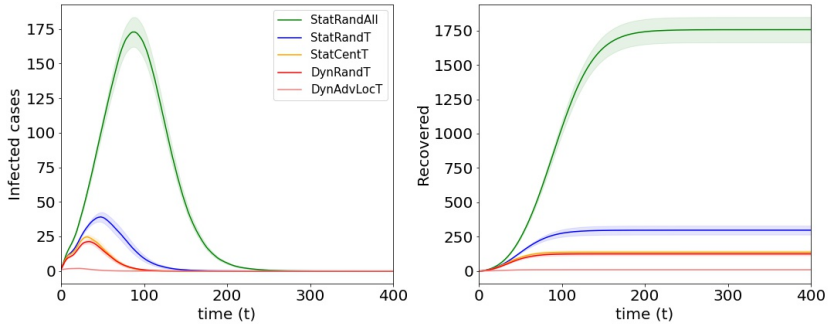

FIGURE A.2: The disease evolution for the proposed campaigns with $\tau = \infty$. For all campaigns, $\mu^- = \mu^+ = 0.001$, the social rate is $\omega = 0.006$, and the target set size is $T = 500$ for static campaigns (i.e., StatRandT and StatCentT) and $T = 50$ for the other dynamic campaigns. For dynamic campaigns, the updating time is $t_r = 20$ for DynRandT, and $t_r = 1$ for DynAdvLocT campaign. Furthermore, for DynAdvLocT, $\zeta = Z = 10$. The results are the average of 50 simulations.

Appendix B

Appendix to Chapter 6

B.1 Degree Distribution for the Real-World Networks

Figure B.1 shows the degree distribution for the real-world networks considered.

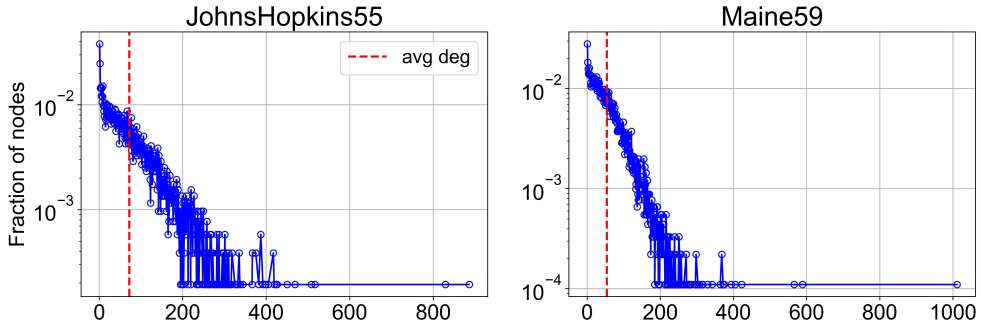

FIGURE B.1: Degree distribution of the considered real-world networks. (A) represents socfb-JohnsHopkins55 network and (B) represents socfb-Maine59 network.

B.2 Dependence of Campaign Efficiency on Budget Size

Figure B.2 shows the results for different network structures with later interventions at $t_s^+ = 200$ for all networks, except for the scale-free network at $t_s^+ = 160$, with approximately 314 ± 14 existing anti-vaccine opinion adopters.

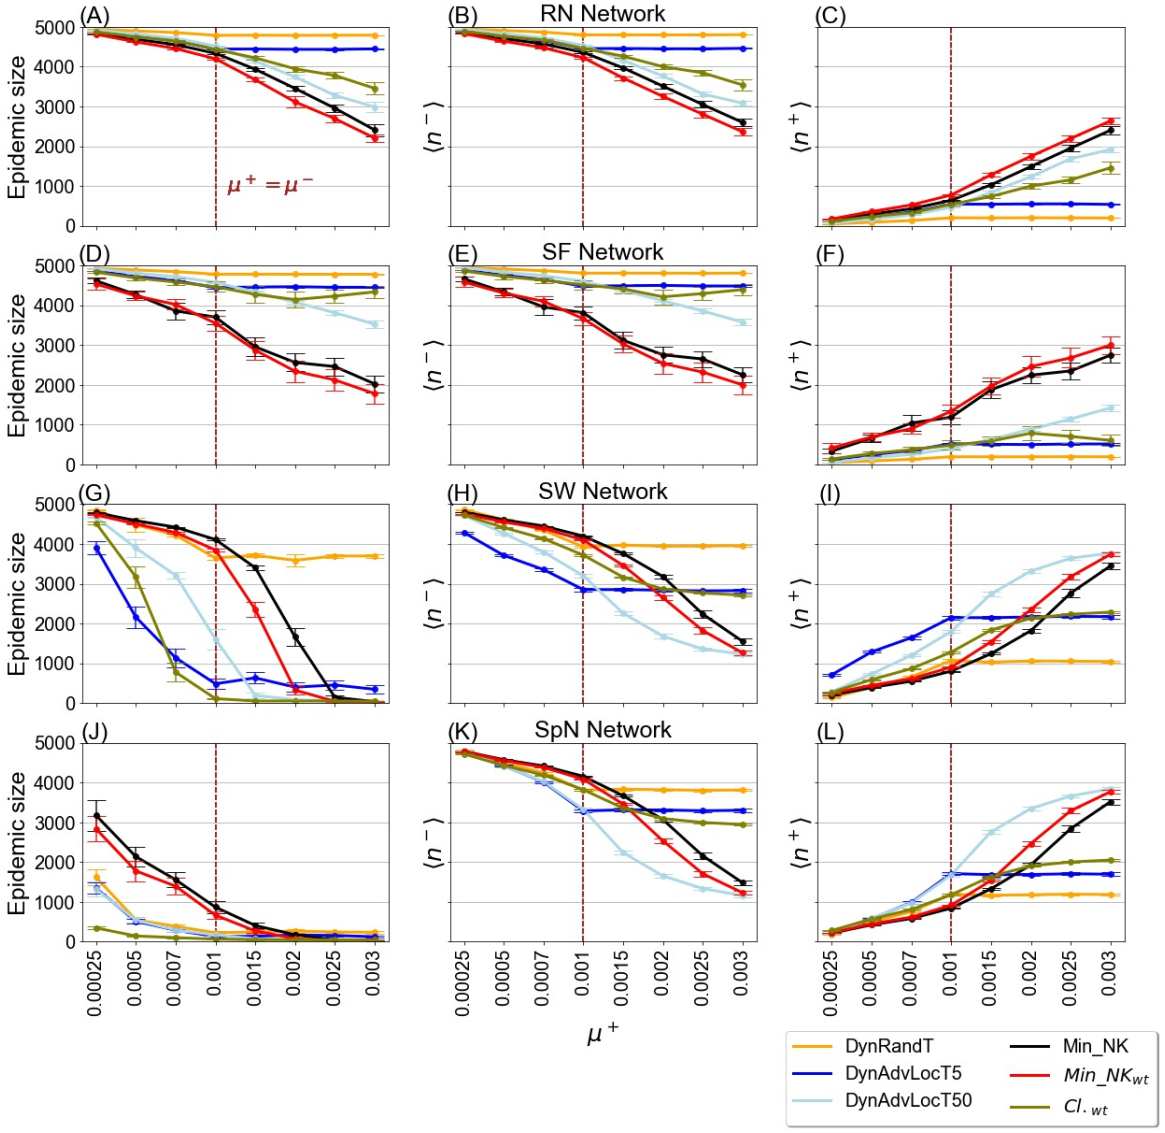

FIGURE B.2: Comparison of the control performance of the proposed intervention campaigns across different network structures as a function of μ^+. Campaigns are applied at $t_s = 200$ in all networks except the scale-free at $t_s = 160$, with an average of 314 ± 14 existing anti-vaccine opinion adopters, approximately 0.1 of the population. With $\mu^- = 0.001$, $\omega = 0.006$, $T = 5$, $t_r = 1$, for DynAdvLocT campaign $\zeta = 1, Z = 10$ for SW and SpN, and $\zeta = 10, Z = 100$ for RN and SF. The results are averaged over 50 simulations.

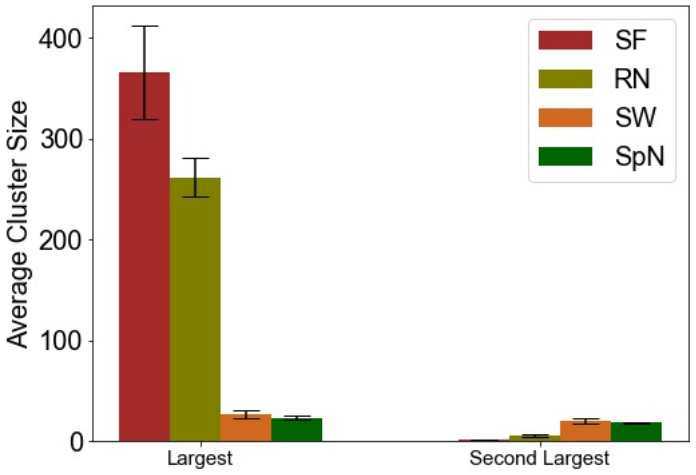

FIGURE B.3: The average size of the largest and second-largest clusters of anti-vaccine communities before applying the positive campaigns at time $t_s^+ = 200$ for all networks, except for scale-free network at $t_s^+ = 160$, with around 0.1 of the population had adopted the anti-vaccine opinion. A total of 20 simulations were conducted. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, and $\omega = 0.006$.

Appendix C

Appendix to Chapter 7

C.1 DynAdvLocT in Multi-Layer Networks

We here presents the epidemic size of the DynAdvLocT campaign as a function of its two main parameters, ζ and Z, which represent the target number of anti-vaccine neighbors and the target number of neutral neighbors, respectively. Our aim is to find the optimal parameter values for this campaigns in multi-layer networks for virtual-based interventions, where information diffusion spreads densely. This follows the approach previously applied to the small-world network in Section 3.4.4.

Figure C.1 shows the final epidemic size as a function of ζ and Z. As shown in the figure, the observed control results align with those seen in the scale-free network (see Fig. 6.1(B)). Specifically, we observe that increasing Z to more than the average degree yields better results, even though the disease spread in different structure (i.e., a small-world network). In scale-free networks, unlike in small-world networks, achieving effective control through DynAdvLocT campaign requires a larger target number of neutrals, Z, than the average degree. This prioritizes nodes with more neutral neighbors than others, typically high-degree nodes in such structures, which provides better control when targeting them[1].

C.2 Uncorrelated Multi-Layer Networks

Figure C.2 illustrates the epidemic size when there is no correlation between the virtual and physical layers, with the virtual layer has a scale-free structure. In this scenario, we also observe that virtual-based interventions outperform physical-based

[1]Note that, the definition of neighbors is the same as in a single-layer network; however, in multi-layer networks, neighbors are defined with respect to the targeted layer, either the virtual layer or the physical layer.

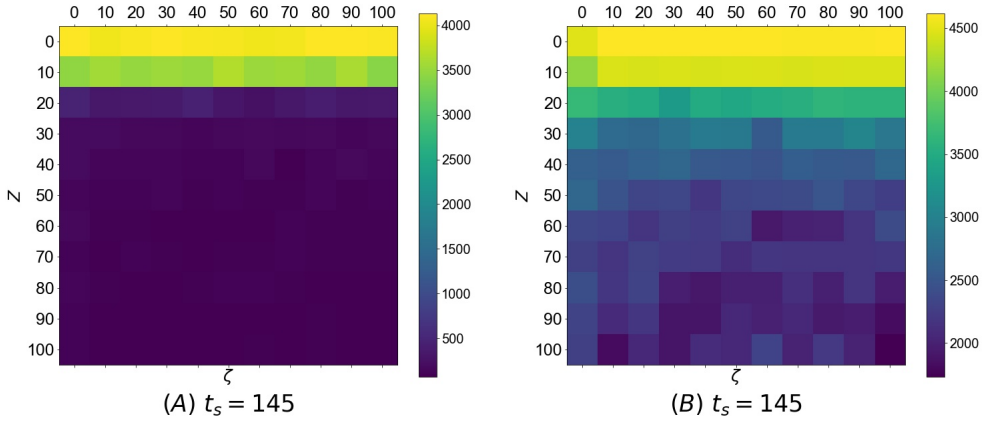

FIGURE C.1: Average epidemic size using DynAdvLocT campaign in multi-layer networks with virtual-based intervention. The average degree $\langle k \rangle$ for $G_v = 14$ and $G_d = 10$. The epidemic size is shown as a function of the target number of anti-vaccine neighbors ζ and the target number of neutral neighbors Z. Applying the positive campaigns based on the information from information layer at time $t_s^+ = 100$ and $t_s^+ = 145$. The experiments are related to a scenario with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, $t_r = 1$, and $T = 5$.

interventions in both early and late intervention settings, consistent with the findings discussed in Section 7.3.3. The campaign performance follows the same trends observed in single-layer scale-free networks, with DynAdvLocT and minimization campaigns (Min_NK and Min_NK$_{wt}$) exhibiting similar control effectiveness.

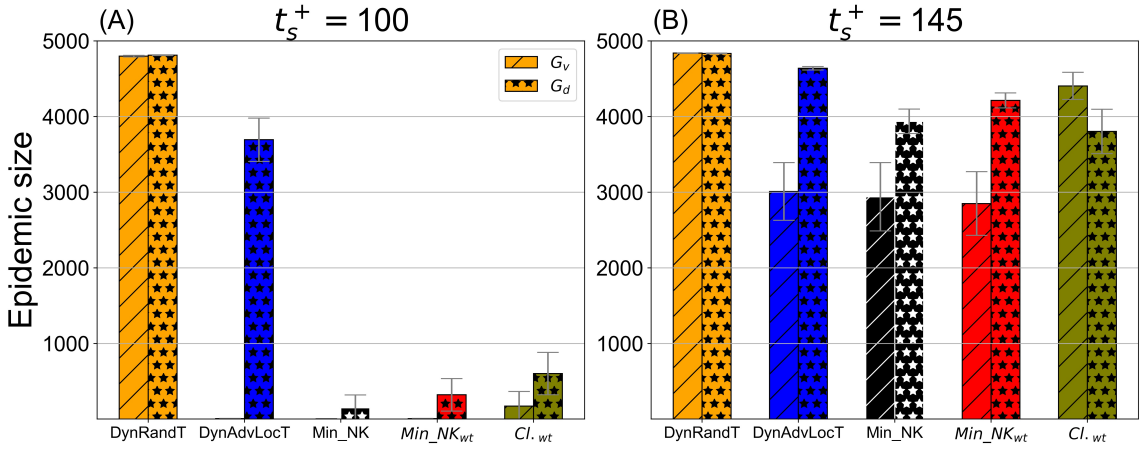

FIGURE C.2: Comparison of the epidemic size resulting from virtual-based (G_v) and physical-based (G_d) intervention campaigns. G_v is a purely scale-free network not extended from the physical layer. The average degree $\langle k \rangle$ for $G_v = 12$ and $G_d = 10$. Intervention campaigns starting at time $t_s^+ = 100$ with the average number of existing anti-vaccine opinion adopters being around 43 $\pm$4 (left panel), and starting at time $t_s^+ = 145$ with the average number of existing anti-vaccine opinion adopters being around 293 $\pm$51 (right panel). Results are the average of 50 realizations with $\mu^+ = \mu^- = 0.001$, $\omega = 0.006$, and $T = 5$.

References

Ten health issues who will tackle this year, 2019. URL https://www.who.int/news-room/spotlight/ten-threats-to-global-health-in-2019.

Measles cases surge worldwide, infecting 10.3 million people in 2023, 2024. URL https://www.who.int/news/item/14-11-2024-measles-cases-surge-worldwide--infecting-10.3-million-people-in-2023.

N Abd Rahim and SM Rafie. Sentiment analysis of social media data in vaccination. *International Journal of Emerging Trends in Engineering Research*, 8(9), 2020.

Daron Acemoglu, Alireza Fallah, Andrea Giometto, Daniel Huttenlocher, Asuman Ozdaglar, Francesca Parise, and Sarath Pattathil. Optimal adaptive testing for epidemic control: combining molecular and serology tests. *Automatica*, 160:111391, 2024.

Gbadebo Collins Adeyanju, Elena Engel, Laura Koch, Tabea Ranzinger, Imtiaz Bin Mohammed Shahid, Micheal G Head, Sarah Eitze, and Cornelia Betsch. Determinants of influenza vaccine hesitancy among pregnant women in europe: a systematic review. *European Journal of Medical Research*, 26:1–12, 2021.

Hadil Alahdal, Fatemah Basingab, and Reem Alotaibi. An analytical study on the awareness, attitude and practice during the covid-19 pandemic in riyadh, saudi arabia. *Journal of Infection and Public Health*, 13(10):1446–1452, 2020.

Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread. In Hocine Cherifi, Murat Donduran, Luis M. Rocha, Chantal Cherifi, and Onur Varol, editors, *Complex Networks & Their Applications XIII*, pages 87–99, Cham, 2025a. Springer Nature Switzerland. ISBN 978-3-031-82439-5.

Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Modelling the mitigation of anti-vaccine opinion propagation to suppress epidemic spread: A computational approach. *PLOS ONE*, 20(3):e0318544, 2025b.

Sarah Alahmadi, Rebecca Hoyle, Michael Head, and Markus Brede. Optimal intervention campaign to combat anti-vaccine social contagion and contain epidemic spread: Impact of network structure. *Applied Network Science*, 10(1):59, 2025c.

Alberto Aleta and Yamir Moreno. Multilayer networks in a nutshell. *Annual Review of Condensed Matter Physics*, 10(1):45–62, 2019.

Aamena Alshamsi, Flávio L Pinheiro, and Cesar A Hidalgo. Optimal diversification strategies in the networks of related products and of related research areas. *Nature Communications*, 9(1):1328, 2018.

Tim Althoff, Pranav Jindal, and Jure Leskovec. Online actions with offline impact: How online social networks influence online and offline user behavior. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 537–546, 2017.

Camilla Ancona, Francesco Lo Iudice, Franco Garofalo, and Pietro De Lellis. A model-based opinion dynamics approach to tackle vaccine hesitancy. *Scientific Reports*, 12(1):11835, 2022.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

Alain Barrat and Martin Weigt. On the properties of small-world network models. *The European Physical Journal B-Condensed Matter and Complex Systems*, 13:547–560, 2000.

Gema Bello-Orgaz, Julio Hernandez-Castro, and David Camacho. Detecting discussion communities on vaccination in twitter. *Future Generation Computer Systems*, 66:125–136, 2017.

Roberto Bernal Jaquez, Luis Angel Alarcón Ramos, and Alexander Schaum. Spreading control in two-layer multiplex networks. *Entropy*, 22(10):1157, 2020.

Ginestra Bianconi. 9the structure of single networks. In *Multilayer Networks: Structure and Function*. Oxford University Press, 06 2018a. ISBN 9780198753919. .

Ginestra Bianconi. Multilayer networks in nature, society and infrastructures. In *Multilayer Networks: Structure and Function*. Oxford University Press, 06 2018b. ISBN 9780198753919. .

Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4-5):175–308, 2006.

Stefano Boccaletti, Ginestra Bianconi, Regino Criado, Charo I Del Genio, Jesús Gómez-Gardenes, Miguel Romance, Irene Sendina-Nadal, Zhen Wang, and Massimiliano Zanin. The structure and dynamics of multilayer networks. *Physics Reports*, 544(1):1–122, 2014.

Ulrik Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.

Markus Brede, Valerio Restocchi, and Sebastian Stein. Transmission errors and influence maximization in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(3):033401, 2019.

Hendrik Bruns, François J Dessart, Michał Krawczyk, Stephan Lewandowsky, Myrto Pantazi, Gordon Pennycook, Philipp Schmid, and Laura Smillie. Investigating the role of source and source trust in prebunks and debunks of misinformation in online experiments across four eu countries. *Scientific Reports*, 14(1):20723, 2024.

Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. Limiting the spread of misinformation in social networks. In *Proceedings of the 20th International Conference on World Wide Web*, pages 665–674, 2011.

Talha Burki. Vaccine misinformation and social media. *The Lancet Digital Health*, 1(6): e258–e259, 2019.

Zhongqi Cai, Markus Brede, and Enrico Gerding. Influence maximization for dynamic allocation in voter dynamics. In *Complex Networks & Their Applications IX: Volume 1, Proceedings of the Ninth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2020*, pages 382–394. Springer, 2021.

Ellsworth Campbell and Marcel Salathé. Complex social contagion makes networks more vulnerable to disease outbreaks. *Scientific Reports*, 3(1):1–6, 2013.

Alessio Cardillo, Catalina Reyes-Suárez, Fernando Naranjo, and Jesús Gómez-Gardenes. Evolutionary vaccination dilemma in complex networks. *Physical Review E*, 88(3):032803, 2013.

Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009.

CDC. Measles cases and outbreaks — data & research. https://www.cdc.gov/measles/data-research/index.html, 2025. Updated December 23, 2025; accessed December 31, 2025.

Centers for Disease Control and Prevention. Measles vaccination, 2025. URL https://www.cdc.gov/measles/vaccines/index.html. Accessed: 2025-05-05.

Damon Centola. The spread of behavior in an online social network experiment. *Science*, 329(5996):1194–1197, 2010.

Damon Centola and Michael Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–734, 2007.

Damon Centola, Víctor M Eguíluz, and Michael W Macy. Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications*, 374(1): 449–456, 2007.

Debayan Chakraborty, Anurag Singh, and Hocine Cherifi. Immunization strategies based on the overlapping nodes in networks with community structure. In *Computational Social Networks: 5th International Conference, CSoNet 2016, Ho Chi Minh City, Vietnam, August 2-4, 2016, Proceedings 5*, pages 62–73. Springer, 2016.

Sukankana Chakraborty, Sebastian Stein, Markus Brede, Ananthram Swami, Geeth de Mel, and Valerio Restocchi. Competitive influence maximisation using voting dynamics. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 978–985, 2019.

Jiangzhuo Chen, Stefan Hoops, Achla Marathe, Henning Mortveit, Bryan Lewis, Srinivasan Venkatramanan, Arash Haddadan, Parantapa Bhattacharya, Abhijin Adiga, Anil Vullikanti, et al. Effective social network-based allocation of covid-19 vaccines. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4675–4683, 2022a.

Jiaxing Chen, Ying Liu, Jing Yue, Xi Duan, and Ming Tang. Coevolving spreading dynamics of negative information and epidemic on multiplex networks. *Nonlinear Dynamics*, pages 1–11, 2022b.

Sijing Chen, Lu Xiao, and Akit Kumar. Spread of misinformation on social media: What contributes to it and how to combat it. *Computers in Human Behavior*, 141: 107643, 2023.

Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 199–208, 2009.

Wei Chen, Laks VS Lakshmanan, and Carlos Castillo. Information and influence propagation in social networks. *Synthesis Lectures on Data Management*, 5(4):1–177, 2013.

Xiaolong Chen, Ruijie Wang, Ming Tang, Shimin Cai, H Eugene Stanley, and Lidia A Braunstein. Suppressing epidemic spreading in multiplex networks with social-support. *New Journal of Physics*, 20(1):013007, 2018.

Tao Cheng, Baoyan Han, and Yunzhe Liu. Exploring public sentiment and vaccination uptake of covid-19 vaccines in england: a spatiotemporal and sociodemographic analysis of twitter data. *Frontiers in Public Health*, 11:1193750, 2023.

Matteo Chinazzi, Jessica T Davis, Marco Ajelli, Corrado Gioannini, Maria Litvinova, Stefano Merler, Ana Pastore y Piontti, Kunpeng Mu, Luca Rossi, Kaiyuan Sun, et al.

The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science*, 368(6489):395–400, 2020.

Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3): 581–588, 1973.

Josh Compton. Inoculation theory. *The SAGE Handbook of Persuasion: Developments in Theory and Practice*, 2:220–236, 2013.

Giulio Corsi. Evaluating twitter's algorithmic amplification of low-credibility content: an observational study. *EPJ Data Science*, 13(1):18, 2024.

Paulo Cesar Ventura da Silva, Fátima Velásquez-Rojas, Colm Connaughton, Federico Vazquez, Yamir Moreno, and Francisco A Rodrigues. Epidemic spreading with awareness and different timescales in multiplex networks. *Physical Review E*, 100(3): 032313, 2019.

Xiangfeng Dai, Peican Zhu, Yangming Guo, and Zhen Wang. Coevolution of vaccination opinions and awareness affecting the spread of epidemics. *IEEE Access*, 7:61558–61569, 2019.

Eleonora D'Andrea, Pietro Ducange, Alessio Bechini, Alessandro Renda, and Francesco Marcelloni. Monitoring the public opinion about the vaccination topic from tweets analysis. *Expert Systems with Applications*, 116:209–226, 2019.

Emily Denniss and Rebecca Lindberg. Social media and the spread of misinformation: infectious and a threat to public health. *Health Promotion International*, 40(2):daaf023, 2025.

Matthew R DeVerna, Francesco Pierri, Yong-Yeol Ahn, Santo Fortunato, Alessandro Flammini, and Filippo Menczer. Modeling the amplification of epidemic spread by individuals exposed to misinformation on social media. *npj Complexity*, 2(1):11, 2025.

Zoltán Dezső and Albert-László Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65(5):055103, 2002.

Lien Anh Ha Do and Kim Mulholland. Measles 2025. *New England Journal of Medicine*, 2025. . URL https://www.nejm.org/doi/full/10.1056/NEJMra2504516.

Yucheng Dong, Min Zhan, Gang Kou, Zhaogang Ding, and Haiming Liang. A survey on the fusion process in opinion dynamics. *Information Fusion*, 43:57–65, 2018.

Claudio O Dorso, Andrés Medus, and Pablo Balenzuela. Vaccination and public trust: A model for the dissemination of vaccination behaviour with external intervention. *Physica A: Statistical Mechanics and its Applications*, 482:433–443, 2017.

Nathan J Doyle, Fergus Cumming, Robin N Thompson, and Michael J Tildesley. When should lockdown be implemented? devising cost-effective strategies for managing epidemics amid vaccine uncertainty. *PLOS Computational Biology*, 20(7):e1012010, 2024.

Eve Dubé, Caroline Laberge, Maryse Guay, Paul Bramadat, Réal Roy, and Julie A Bettinger. Vaccine hesitancy: an overview. *Human Vaccines & Immunotherapeutics*, 9 (8):1763–1773, 2013.

Eve Dubé, Maryline Vivion, and Noni E MacDonald. Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. *Expert Review of Vaccines*, 14(1):99–117, 2015.

P Erdős and A Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6 (290-297):18, 1959.

Şirag Erkol, Claudio Castellano, and Filippo Radicchi. Systematic comparison between methods for the detection of influential spreaders in complex networks. *Scientific Reports*, 9(1):15095, 2019.

Soheil Eshghi, Victor M Preciado, Saswati Sarkar, Santosh S Venkatesh, Qing Zhao, Raissa D'Souza, and Ananthram Swami. Spread, then target, and advertise in waves: optimal budget allocation across advertising channels. *IEEE Transactions on Network Science and Engineering*, 7(2):750–763, 2018.

W Douglas Evans, Jeffrey B Bingenheimer, Michael Long, Khadidiatou Ndiaye, Dante Donati, Nandan M Rao, Selinam Akaba, Ifeanyi Nsofor, and Sohail Agha. Outcomes of a social media campaign to promote covid-19 vaccination in nigeria. *PLOS ONE*, 18(9):e0290757, 2023.

Salvatore Federico, Giorgio Ferrari, and Maria-Laura Torrente. Optimal vaccination in a sirs epidemic model. *Economic Theory*, 77(1):49–74, 2024.

Feng Fu, Daniel I Rosenbloom, Long Wang, and Martin A Nowak. Imitation dynamics of vaccination behaviour on social networks. *Proceedings of the Royal Society B: Biological Sciences*, 278(1702):42–49, 2011.

Marzena Fügenschuh and Feng Fu. Overcoming vaccine hesitancy by multiplex social network targeting: an analysis of targeting algorithms and implications. *Applied Network Science*, 8(1):67, 2023.

Sebastian Funk, Erez Gilad, Chris Watkins, and Vincent AA Jansen. The spread of awareness and its impact on epidemic outbreaks. *Proceedings of the National Academy of Sciences*, 106(16):6872–6877, 2009.

Ashley Gromis and Ka-Yuet Liu. Spatial clustering of vaccine exemptions on the risk of a measles outbreak. *Pediatrics*, 149(1), 2022.

Adrien Guille, Hakim Hacid, Cecile Favre, and Djamel A Zighed. Information diffusion in online social networks: A survey. *ACM Sigmod Record*, 42(2):17–28, 2013.

Naveen Gupta, Anurag Singh, and Hocine Cherifi. Community-based immunization strategies for epidemic control. In *2015 7th International Conference on Communication Systems and Networks (COMSNETS)*, pages 1–6. IEEE, 2015.

GS Harari and LHA Monteiro. An epidemic model with pro and anti-vaccine groups. *Acta Biotheoretica*, 70(3):20, 2022.

Oliver Hein, Michael Schwind, and Wolfgang König. Scale-free networks: The impact of fat tailed degree distribution on diffusion and communication processes. *Wirtschaftsinformatik*, 48:267–275, 2006.

Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4): 599–653, 2000.

Peter J Hotez, Tasmiah Nuzhath, and Brian Colwell. Combating vaccine hesitancy and other 21st century social determinants in the global fight against measles. *Current Opinion in Virology*, 41:1–7, 2020.

Thomas House, Joshua V Ross, and David Sirl. How big is an outbreak likely to be? methods for epidemic final-size calculation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 469(2150):20120436, 2013.

Tao Hu, Siqin Wang, Wei Luo, Mengxi Zhang, Xiao Huang, Yingwei Yan, Regina Liu, Kelly Ly, Viraj Kacker, Bing She, et al. Revealing public opinion towards covid-19 vaccines with twitter data in the united states: spatiotemporal perspective. *Journal of Medical Internet Research*, 23(9):e30854, 2021.

He Huang, Yahong Chen, and Yefeng Ma. Modeling the competitive diffusions of rumor and knowledge and the impacts on epidemic spreading. *Applied Mathematics and Computation*, 388:125536, 2021.

Elizabeth Hunter, Brian Mac Namee, and John D. Kelleher. A comparison of agent-based models and equation-based models for infectious disease epidemiology. In *26th AIAI Irish Conference on Artificial Intelligence and Cognitive Science*, 2018. .

Genki Ichinose and Takehiro Kurisaku. Positive and negative effects of social impact on evolutionary vaccination game in networks. *Physica A: Statistical Mechanics and its Applications*, 468:84–90, 2017.

Jian Jiang and Tianshou Zhou. Resource control of epidemic spreading through a multilayer network. *Scientific Reports*, 8(1):1629, 2018.

KM Ariful Kabir and Jun Tanimoto. Vaccination strategies in a two-layer sir/v–ua epidemic model with costly information and buzz effect. *Communications in Nonlinear Science and Numerical Simulation*, 76:92–108, 2019.

KM Ariful Kabir, Kazuki Kuga, and Jun Tanimoto. Effect of information spreading to suppress the disease contagion on the epidemic vaccination game. *Chaos, Solitons & Fractals*, 119:180–187, 2019.

Kundan Kandhway and Joy Kuri. How to run a campaign: Optimal control of sis and sir information epidemics. *Applied Mathematics and Computation*, 231:79–92, 2014.

Tapan Kumar Kar and Soovoojeet Jana. A theoretical study on mathematical modelling of an infectious disease with application of optimal control. *Biosystems*, 111(1):37–50, 2013.

Matt J Keeling and Ken TD Eames. Networks and epidemic models. *Journal of The Royal Society Interface*, 2(4):295–307, 2005.

Matt J. Keeling and Pejman Rohani. *Introduction to Simple Epidemic Models*, pages 15–53. Princeton University Press, 2008. ISBN 9780691116174. URL http://www.jstor.org/stable/j.ctvcm4gk0.5.

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

William Ogilvy Kermack and Anderson G McKendrick. 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.

Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 163–170, 2000.

Emmanuel F Koku, Nettie Johnson-Yengbeh, and Ava Muhr. Addressing covid-19 vaccine hesitancy and uptake among african immigrants: Lessons from a community-based outreach program. *Journal of Racial and Ethnic Health Disparities*, pages 1124–1138, 2025.

Pinelopi Konstantinou, Katerina Georgiou, Navin Kumar, Maria Kyprianidou, Christos Nicolaides, Maria Karekla, and Angelos P Kassianos. Transmission of vaccination attitudes and uptake based on social contagion theory: A scoping review. *Vaccines*, 9(6):607, 2021.

Vito Latora and Massimo Marchiori. Efficient behavior of small-world networks. *Physical Review Letters*, 87(19):198701, 2001.

Vito Latora, Vincenzo Nicosia, and Giovanni Russo. *Centrality Measures*, page 31–68. Cambridge University Press, 2017. .

Sun Kyong Lee, Juhyung Sun, Seulki Jang, and Shane Connelly. Misinformation of covid-19 vaccines and vaccine hesitancy. *Scientific Reports*, 12(1):13681, 2022.

Thomas Lemaitre, Nathalie Carrier, Anne Farrands, Virginie Gosselin, Geneviève Petit, and Arnaud Gagneur. Impact of a vaccination promotion intervention using motivational interview techniques on long-term vaccine coverage: the promovac strategy. *Human Vaccines & Immunotherapeutics*, 15(3):732–739, 2019.

Hao Li, Luqi Wang, Mengxi Zhang, Yihan Lu, and Weibing Wang. Effects of vaccination and non-pharmaceutical interventions and their lag times on the covid-19 pandemic: Comparison of eight countries. *PLOS Neglected Tropical Diseases*, 16(1):e0010101, 2022.

Tracy A Lieu, G Thomas Ray, Nicola P Klein, Cindy Chung, and Martin Kulldorff. Geographic clusters in underimmunization and vaccine refusal. *Pediatrics*, 135(2): 280–289, 2015.

Weiyi Liu, Kun Yue, Hong Wu, Jin Li, Donghua Liu, and Duanping Tang. Containment of competitive influence spread in social networks. *Knowledge-Based Systems*, 109:266–275, 2016.

Simon Lynch, Long Tran-Thanh, Francisco C Santos, et al. Fostering cooperation in structured populations through local and global interference strategies. In *27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence*, pages 289–295. IJCAI, 2018.

Christopher Lynn and Daniel D Lee. Maximizing influence in an ising network: A mean-field optimal solution. *Advances in Neural Information Processing Systems*, 29, 2016.

Ziyu Ma and Elena Parilina. Stable multiplex networks: definitions and characterizations. *Journal of Complex Networks*, 12(5):cnae040, 2024.

Roy Rillera Marzo, Absar Ahmad, Md Saiful Islam, Mohammad Yasir Essar, Petra Heidler, Isabel King, Arulmani Thiyagarajan, Kittisak Jermsittiparsert, Karnjana Songwathana, Delan Ameen Younus, et al. Perceived covid-19 vaccine effectiveness, acceptance, and drivers of vaccination decision-making among the general adult population: A global survey of 20 countries. *PLOS Neglected Tropical Diseases*, 16(1): e0010103, 2022.

Antonia Maria Masucci and Alonso Silva. Strategic resource allocation for competitive influence in social networks. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 951–958. IEEE, 2014.

Naoki Masuda. Opinion control in complex networks. *New Journal of Physics*, 17(3): 033031, 2015.

Jennifer L Matas, Latrice G Landry, LaTasha Lee, Shantoy Hansel, Makella S Coudray, Lina V Mata-McMurry, Nishanth Chalasani, Liou Xu, Taylor Stair, Christina Edwards, et al. Demographic determinants and geographical variability of covid-19 vaccine hesitancy in underserved communities: cross-sectional study. *JMIR Public Health and Surveillance*, 9(1):e34163, 2023.

Laura Matrajt, Julia Eaton, Tiffany Leung, and Elizabeth R Brown. Vaccine optimization for covid-19: Who to vaccinate first? *Science Advances*, 7(6):eabf1374, 2021.

Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.

Rohan S Mehta and Noah A Rosenberg. Modelling anti-vaccine sentiment as a cultural pathogen. *Evolutionary Human Sciences*, 2, 2020.

Paul Mena. Cleaning up social media: The effect of warning labels on likelihood of sharing false news on facebook. *Policy & Internet*, 12(2):165–183, 2020.

Xueyu Meng, Sijie Han, Leilei Wu, Shubin Si, and Zhiqiang Cai. Analysis of epidemic vaccination strategies by node importance and evolutionary game on complex networks. *Reliability Engineering & System Safety*, 219:108256, 2022.

Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, pages 29–42, 2007.

Jayanta Mondal and Subhas Khajanchi. Mathematical modeling and optimal intervention strategies of the covid-19 outbreak. *Nonlinear Dynamics*, 109(1):177–202, 2022.

Bjarke Mønsted and Sune Lehmann. Characterizing polarization in online vaccine discourse—a large-scale study. *PLOS ONE*, 17(2):e0263746, 2022.

Felipe Montes, Ana María Jaramillo, Jose D Meisel, Albert Diaz-Guilera, Juan A Valdivia, Olga L Sarmiento, and Roberto Zarama. Benchmarking seeding strategies for spreading processes in social networks: an interplay between influencers, topologies and sizes. *Scientific Reports*, 10(1):3666, 2020.

Danielle Xiaodan Morales, Tyler Fox Beltran, and Stephanie Alexandra Morales. Gender, socioeconomic status, and covid-19 vaccine hesitancy in the us: an intersectionality approach. *Sociology of Health & Illness*, 44(6):953–971, 2022.

Nabeela Mumtaz, Caroline Green, and Jim Duggan. Exploring the effect of misinformation on infectious disease transmission. *Systems*, 10(2):50, 2022.

Arvind Narayanan. Understanding social media recommendation algorithms. 2023.

Mark Newman. *Networks*. Oxford University Press, 07 2018. ISBN 9780198805090. .

Mark EJ Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1): 016128, 2002.

Veronija Nolcheska. The influence of social networks on consumer behavior. *Cataloging-In-Publication Data*, 95, 2017.

Cameron Nowzari, Victor M Preciado, and George J Pappas. Optimal resource allocation for control of networked epidemic models. *IEEE Transactions on Control of Network Systems*, 4(2):159–169, 2015.

Saad B Omer, Daniel A Salmon, Walter A Orenstein, M Patricia Dehart, and Neal Halsey. Vaccine refusal, mandatory immunization, and the risks of vaccine-preventable diseases. *New England Journal of Medicine*, 360(19):1981–1988, 2009.

Juliette Paireau, Marie-Laure Charpignon, Sophie Larrieu, Clémentine Calba, Nathanaël Hozé, Pierre-Yves Boëlle, Rodolphe Thiebaut, Mélanie Prague, and Simon Cauchemez. Impact of non-pharmaceutical interventions, weather, vaccination, and variants on covid-19 transmission across departments in france. *BMC Infectious Diseases*, 23(1):190, 2023.

Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of Modern Physics*, 87 (3):925–979, 2015.

María Pérez-Ortiz, Petru Manescu, Fabio Caccioli, Delmiro Fernández-Reyes, Parashkev Nachev, and John Shawe-Taylor. Network topological determinants of pathogen spread. *Scientific Reports*, 12(1):7692, 2022.

Dung V Pham, Giang L Nguyen, Tu N Nguyen, Canh V Pham, and Anh V Nguyen. Multi-topic misinformation blocking with budget constraint on online social networks. *IEEE Access*, 8:78879–78889, 2020.

Francesco Pierri, Brea L Perry, Matthew R DeVerna, Kai-Cheng Yang, Alessandro Flammini, Filippo Menczer, and John Bryden. Online misinformation is linked to early covid-19 vaccination hesitancy and refusal. *Scientific Reports*, 12(1):5966, 2022.

William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. Section 10.10. linear programming: The simplex method. In *Numerical Recipes: The Art of Scientific Computing*, 2007.

Rafael Prieto Curiel and Humberto González Ramírez. Vaccination strategies against covid-19 and the diffusion of anti-vaccination views. *Scientific Reports*, 11(1):6626, 2021.

Neha Puri, Eric A Coomes, Hourmazd Haghbayan, and Keith Gunaratne. Social media and vaccine hesitancy: new updates for the era of covid-19 and globalized infectious diseases. *Human Vaccines & Immunotherapeutics*, 16(11):2586–2593, 2020.

Geofakta Razali, Masfiatun Nikmah, I Nyoman Tri Sutaguna, PA Andiena Nindya Putri, and Muhammad Yusuf. The influence of viral marketing and social media marketing on instagram adds purchase decisions., 2023.

Jonathan M Read, Ken TD Eames, and W John Edmunds. Dynamic social networks and the implications for the spread of infectious disease. *Journal of The Royal Society Interface*, 5(26):1001–1007, 2008.

Stephanie M Reich, Kaveri Subrahmanyam, and Guadalupe Espinoza. Friending, iming, and hanging out face-to-face: overlap in adolescents' online and offline social networks. *Developmental Psychology*, 48(2):356, 2012.

Julie Riddell, Srebrenka Letina, Kathryn Skivington, Daniel Archambault, Valerie Wells, Emily Long, Ruth Hunter, Mark McCann, et al. Methods for interventions using networks to improve health: A narrative synthesis of methodological research on network data collection, visualisation and intervention. *Social Networks*, 2025.

Arne Roets et al. 'fake news': Incorrect, but hard to correct. the role of cognitive ability on the impact of false information on social impressions. *Intelligence*, 65:107–110, 2017.

Daniel M Romero, Brendan Meeder, and Jon Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th International Conference on World Wide Web*, pages 695–704, 2011.

Guillermo Romero Moreno, Long Tran-Thanh, and Markus Brede. Continuous influence maximisation for the voter dynamics: Is targeting high-degree nodes a good strategy? In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, page 1981–1983, Richland, SC, 2020. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450375184.

Guillermo Romero Moreno, Sukankana Chakraborty, and Markus Brede. Shadowing and shielding: Effective heuristics for continuous influence maximisation in the voting dynamics. *PLOS ONE*, 16(6):e0252515, 2021.

Jon Roozenbeek, Sander Van Der Linden, and Thomas Nygren. Prebunking interventions based on "inoculation" theory can reduce susceptibility to misinformation across cultures. *Harvard Kennedy School Misinformation Review*, 2020.

Jon Roozenbeek, Sander Van Der Linden, Beth Goldberg, Steve Rathje, and Stephan Lewandowsky. Psychological inoculation improves resilience against misinformation on social media. *Science Advances*, 8(34):eabo6254, 2022.

Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. www.networkrepository.com.

Zhongyuan Ruan, Ming Tang, and Zonghua Liu. Epidemic spreading with information-driven vaccination. *Physical Review E*, 86(3):036117, 2012.

Eldar Sadikov and Maria Montserrat Medina Martinez. Information propagation on twitter. *CS322 project report*, 23, 2009.

Anwar Said, Timothy D Bowman, Rabeeh Ayaz Abbasi, Naif Radi Aljohani, Saeed-Ul Hassan, and Raheel Nawaz. Mining network-level properties of twitter altmetrics data. *Scientometrics*, 120(1):217–235, 2019.

Marcel Salathé and Sebastian Bonhoeffer. The effect of opinion clustering on disease outbreaks. *Journal of The Royal Society Interface*, 5(29):1505–1508, 2008.

Marcel Salathé and James H Jones. Dynamics and control of diseases in networks with community structure. *PLOS Computational Biology*, 6(4):e1000736, 2010.

Marcel Salathé and Shashank Khandelwal. Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control. *PLOS Computational Biology*, 7(10):e1002199, 2011.

Ana Lucía Schmidt, Fabiana Zollo, Antonio Scala, Cornelia Betsch, and Walter Quattrociocchi. Polarization of the vaccination debate on facebook. *Vaccine*, 36(25): 3606–3612, 2018.

Haeseung Seo, Aiping Xiong, Sian Lee, and Dongwon Lee. If you have a reliable source, say something: effects of correction comments on covid-19 misinformation. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 896–907, 2022.

Liuyan Shi, Liang Zhang, and Yun Lu. Evaluating social network-based weight loss interventions in chinese population: An agent-based simulation. *PLOS ONE*, 15(8): e0236716, 2020.

Qihao Shi, Can Wang, Deshi Ye, Jiawei Chen, Yan Feng, and Chun Chen. Adaptive influence blocking: Minimizing the negative spread by observation-based policies. In *2019 IEEE 35th international Conference on Data Engineering (ICDE)*, pages 1502–1513. IEEE, 2019.

Donghee Shin and Emily Y Shin. Cascading falsehoods: mapping the diffusion of misinformation in algorithmic environments. *AI & SOCIETY*, pages 1–18, 2025.

Alina Sîrbu, Vittorio Loreto, Vito DP Servedio, and Francesca Tria. Opinion dynamics: models, extensions and external effects. In *Participatory Sensing, Opinions and Collective Awareness*, pages 363–401. Springer International Publishing, 2017.

Andrei Sontag, Tim Rogers, and Christian A Yates. Misinformation can prevent the suppression of epidemics. *Journal of the Royal Society Interface*, 19(188):20210668, 2022.

Vishal Sood and Sidney Redner. Voter model on heterogeneous graphs. *Physical Review Letters*, 94(17):178701, 2005.

Lihong Sun, Qiang He, Yueyang Teng, Qi Zhao, Xin Yan, and Xingwei Wang. A complex network-based vaccination strategy for infectious diseases. *Applied Soft Computing*, 136:110081, 2023.

Yuichi Tatsukawa, Md Rajib Arefin, Masaki Tanaka, and Jun Tanimoto. Free ticket, discount ticket or intermediate of the best of two worlds–which subsidy policy is socially optimal to suppress the disease spreading? *Journal of Theoretical Biology*, 520: 110682, 2021.

Li Qian Tay, Mark J Hurlstone, Tim Kurz, and Ullrich KH Ecker. A comparison of prebunking and debunking interventions for implied versus explicit misinformation. *British Journal of Psychology*, 113(3):591–607, 2022.

Emma Teasdale, Miriam Santer, Adam WA Geraghty, Paul Little, and Lucy Yardley. Public perceptions of non-pharmaceutical interventions for reducing transmission of respiratory infection: systematic review and synthesis of qualitative studies. *BMC Public Health*, 14(1):589, 2014.

Qawi K Telesford, Karen E Joyce, Satoru Hayasaka, Jonathan H Burdette, and Paul J Laurienti. The ubiquity of small-world networks. *Brain Connectivity*, 1(5):367–375, 2011.

Sarah Thompson, Johanna C Meyer, Rosemary J Burnett, and Stephen M Campbell. Mitigating vaccine hesitancy and building trust to prevent future measles outbreaks in england. *Vaccines*, 11(2):288, 2023.

Amo Tong, Ding-Zhu Du, and Weili Wu. On misinformation containment in online social networks. *Advances in Neural Information Processing Systems*, 31, 2018.

Sho Tsugawa and Hiroyuki Ohsaki. Negative messages spread rapidly and widely on social media. In *Proceedings of the 2015 ACM on Conference on Online Social Networks*, pages 151–160, 2015.

Areeba Umair and Elio Masciari. Sentimental and spatial analysis of covid-19 vaccines tweets. *Journal of Intelligent Information Systems*, 60(1):1–21, 2023.

Imre Varga. Comparison of network topologies by simulation of advertising. In *International Conference on Complexity, Future Information Systems and Risk*, volume 2, pages 17–22. SCITEPRESS, 2017.

Patrick Vinck, Phuong N Pham, Kenedy K Bindu, Juliet Bedford, and Eric J Nilles. Institutional trust and misinformation in the response to the 2018–19 ebola outbreak in north kivu, dr congo: a population-based survey. *The Lancet Infectious Diseases*, 19 (5):529–536, 2019.

Russell M Viner, Simon J Russell, Helen Croker, Jessica Packer, Joseph Ward, Claire Stansfield, Oliver Mytton, Chris Bonell, and Robert Booy. School closure and management practices during coronavirus outbreaks including covid-19: a rapid systematic review. *The Lancet Child & Adolescent Health*, 4(5):397–404, 2020.

Maryline Vivion, Elhadji Anassour Laouan Sidi, Cornelia Betsch, Maude Dionne, Eve Dubé, S Michelle Driedger, Dominique Gagnon, Janice Graham, Devon Greyson, Denis Hamel, et al. Prebunking messaging to inoculate against covid-19 vaccine misinformation: an effective strategy for public health. *Journal of Communication in Healthcare*, 15(3):232–242, 2022.

MyVan Vo, Joshua A Glasser, and Zhilan Feng. Optimal allocation of resources to healthcare workers or the general populace: a modelling study. *Royal Society Open Science*, 8(11):210823, 2021.

Emily K Vraga and Leticia Bode. I do not believe you: How providing a source corrects health misperceptions across social media platforms. *Information, Communication & Society*, 21(10):1337–1353, 2018.

Hongjian Wang and Yajia Lan. The global dynamic transmissibility of covid-19 and its influencing factors: an analysis of control measures from 176 countries. *BMC Public Health*, 23(1):404, 2023.

Xinyu Wang, Danyang Jia, Shupeng Gao, Chengyi Xia, Xuelong Li, and Zhen Wang. Vaccination behavior by coupling the epidemic spreading with the human decision under the game theory. *Applied Mathematics and Computation*, 380:125232, 2020.

Zhen Wang, Chris T Bauch, Samit Bhattacharyya, Alberto d'Onofrio, Piero Manfredi, Matjaž Perc, Nicola Perra, Marcel Salathé, and Dawei Zhao. Statistical physics of vaccination. *Physics Reports*, 664:1–113, 2016.

Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

Arie Wahyu Wijayanto and Tsuyoshi Murata. Effective and scalable methods for graph protection strategies against epidemics on dynamic networks. *Applied Network Science*, 4(1):1–31, 2019.

World Health Organization. Infodemic. URL https://www.who.int/health-topics/infodemic#tab=tab_1. Accessed: 2025-05-22.

Jiwei Xu, Jincheng Li, Zhen Han, and Peican Zhu. Coupled epidemic dynamics with awareness heterogeneity in multiplex networks. *Chaos, Solitons & Fractals*, 187: 115335, 2024.

Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, volume 16, pages 740–748, 2016.

Lan Yang, Zhiwu Li, and Alessandro Giua. Rumor containment by spreading correct information in social networks. In *2019 American Control Conference (ACC)*, pages 5608–5613. IEEE, 2019a.

Lan Yang, Zhiwu Li, and Alessandro Giua. Containment of rumor spread in complex social networks. *Information Sciences*, 506:113–130, 2020.

Yingrui Yang, Ashley McKhann, Sixing Chen, Guy Harling, and Jukka-Pekka Onnela. Efficient vaccination strategies for epidemic control using network information. *Epidemics*, 27:115–122, 2019b.

Qian Yin, Zhishuang Wang, Chengyi Xia, and Chris T Bauch. Impact of co-evolution of negative vaccine-related information, vaccination behavior and epidemic spreading in multilayer networks. *Communications in Nonlinear Science and Numerical Simulation*, 109:106312, 2022.

Qian Yin, Zhishuang Wang, and Chengyi Xia. Information-epidemic co-evolution propagation under policy intervention in multiplex networks. *Nonlinear Dynamics*, 111(15):14583–14595, 2023.

Samira Yousefinaghani, Rozita Dara, Samira Mubareka, Andrew Papadopoulos, and Shayan Sharif. An analysis of covid-19 vaccine sentiments and opinions on twitter. *International Journal of Infectious Diseases*, 108:256–262, 2021.

Xiaoyi Yuan, Ross J Schuchard, and Andrew T Crooks. Examining emergent communities and social bots within the polarized online vaccination debate in twitter. *Social Media+ Society*, 5(3):2056305119865465, 2019.

Ahmad Zareie and Rizos Sakellariou. Minimizing the spread of misinformation in online social networks: A survey. *Journal of Network and Computer Applications*, 186: 103094, 2021.

Ahmad Zareie and Rizos Sakellariou. Rumour spread minimization in social networks: A source-ignorant approach. *Online Social Networks and Media*, 29:100206, 2022.

Hai-Feng Zhang, Zhi-Xi Wu, Ming Tang, and Ying-Cheng Lai. Effects of behavioral response and vaccination policy on epidemic spreading-an approach based on evolutionary-game dynamics. *Scientific Reports*, 4(1):1–10, 2014.

Hai-Feng Zhang, Pan-Pan Shu, Zhen Wang, Ming Tang, and Michael Small. Preferential imitation can invalidate targeted subsidy policies on seasonal-influenza diseases. *Applied Mathematics and Computation*, 294:332–342, 2017.

Haifeng Zhang, Jie Zhang, Changsong Zhou, Michael Small, and Binghong Wang. Hub nodes inhibit the outbreak of epidemic under voluntary vaccination. *New Journal of Physics*, 12(2):023015, 2010.

Huiyuan Zhang, Huiling Zhang, Xiang Li, and My T Thai. Limiting the spread of misinformation while effectively raising awareness in social networks. In *International Conference on Computational Social Networks*, pages 35–47. Springer, 2015.

Junlong Zhang and Yu Luo. Degree centrality, betweenness centrality, and closeness centrality in social network. In *Proceedings of the 2017 2nd international Conference on Modelling, Simulation and Applied Mathematics (MSAM2017)*, pages 300–303. Atlantis press, 2017.

Wenyi Zhang, Yao Wu, Bo Wen, Yongming Zhang, Yong Wang, Wenwu Yin, Shanhua Sun, Xianyu Wei, Hailong Sun, Zhijie Zhang, et al. Non-pharmaceutical interventions for covid-19 reduced the incidence of infectious diseases: a controlled interrupted time-series study. *Infectious Diseases of Poverty*, 12(02):60–71, 2023.