Heterogeneity and homophily in coupled behavior-disease dynamics: from model structure to early warnings

by

Zitao He

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2025

© Zitao He 2025

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

Supervisor:	Chris T. Bauch Professor, Department of Applied Mathematics, University of Waterloo
Internal Members:	Roberto Guglielmi Assistant Professor, Department of Applied Mathematics, University of Waterloo
	N. Sri Namachchivaya Professor, Department of Applied Mathematics, University of Waterloo
Internal-External Member:	Audrey Béliveau Associate Professor, Department of Statistics and Actuarial Science, University of Waterloo
External Examiner:	Jacques Bélair Professor, Department of Mathematics and Statistics, Université de Montréal

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

- Chapter 2: The work in this chapter was performed by Zitao He (ZH) under the supervision of Dr. Chris T. Bauch (CTB), resulting in a co-authored paper published in *Mathematical Biosciences* [109]. Contributions were as follows: CTB conceived of the study. ZH curated the data, analyzed the model, and wrote the original draft. All authors developed the model and revised the manuscript.

- Chapter 3: The work in this chapter was performed by ZH under the supervision of CTB, resulting in a co-authored paper published in *Mathematical Biosciences and Engineering* [110]. Contributions were as follows: CTB conceived of the study. ZH curated the data, developed the model, analyzed the model, and wrote the original draft. All authors revised the manuscript.

- Chapter 4: The work in this chapter was performed by ZH under the supervision of CTB, in collaboration with Dr. Piero Manfredi (PM, University of Pisa) and Dr. Alberto d'Onofrio (AdO, University of Trieste), resulting in a co-authored manuscript under review at *Proceedings of the Royal Society B: Biological Sciences*. Contributions were as follows: ZH conceived of the study, developed the model, analyzed the model, and wrote the original draft. All authors revised the manuscript.

Abstract

Understanding how human behavior and infectious disease dynamics interact is essential for anticipating and mitigating outbreaks. While coupled behavior-disease models have provided valuable insights into the feedback between disease transmission and vaccination behavior, many assume homogeneous populations and neglect the influence of social structure in shaping individual vaccination strategies. Traditional surveillance systems often lack timely data on vaccination behavior, making it difficult to monitor changes in public vaccine sentiment. Moreover, existing statistical methods for detecting early warning signals of critical transitions rely on assumptions that do not always hold in real-world settings. This thesis addresses these limitations by incorporating population heterogeneity and homophily into a coupled behavior-disease model, and by using the resulting simulations to support the training of data-driven models for forecasting outbreak risks from high-frequency social media data. Specifically, we develop a coupled behavior-disease model that distinguishes social media users from non-users, capturing indirect heterogeneity in how individuals access vaccine-related information. The model demonstrates that homophily slows the spread of pro-vaccine strategies, pushing the population closer to tipping points. It also suggests that early vaccine-related online discussions may offer predictive signals of future outbreaks. Building on these findings, we generate synthetic time series with heavy-tailed noise to mimic real-world social media data. These model-generated data are used to train deep learning classifiers, under CNN-LSTM and ResNet architectures, to detect early warning signals in social media data. These classifiers outperform conventional statistical indicators, such as variance and lag-1 autocorrelation, in both sensitivity and specificity. Finally, we extend the modeling framework to a generalized multi-group vaccination game, considering direct heterogeneity in levels of vaccine support. Simulations reveal that homophily contributes to the persistence of opinion polarization in the population, regardless of the presence of diseases. Together, these studies highlight the need to account for heterogeneity in modeling vaccination behavior and that homophily can have various effects depending on the states of the system. We also show that combining mechanistic models and data-driven techniques can help detect emerging risks of disease outbreaks, informing more proactive public health policies.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor, Dr. Chris Bauch, for the trust and freedom you gave me to explore the research topics I am most passionate about. Your guidance and support throughout this five-year journey have been invaluable–whenever I hit a roadblock, your knowledge and experience consistently helped me find direction and move forward. I would like to thank my examining committee members, Dr. Roberto Guglielmi, Dr. Sri Namachchivaya, Dr. Audrey Béliveau, and Dr. Jacques Bélair, for dedicating their time to read and evaluate this thesis.

I would like to thank Dr. Piero Manfredi and Dr. Alberto d'Onofrio for their insightful contributions to the fourth chapter of this thesis. Thanks to Dr. Iain Moyles and Dr. Abba Gumel for giving me opportunities to present my research at CAIMS 2024 and SMB 2025, and to Dr. Claus Kadelka for inviting me to give a seminar talk at Iowa State University on my very first publication. This kind of recognition is encouraging for a young researcher.

I would like to thank two important mentors I met during my upper years at the University of Toronto, Dr. Shenhui Liu and Federico Cao, who inspired and encouraged me to pursue research and solidified my decision to undertake a doctoral degree in applied mathematics.

To my lab mates, Amrita, Athira, Russell, Mackenzie, Yazdan, Sefah, Ali, Shijie, and our de facto lab members, Maliha and Megha: thank you for the office chats, lunch breaks, research discussions, and all the good memories. I truly enjoyed the time I spent in the lab; it was always worth the long trip from home. Also, I would like to thank my fellow students in the department, Albert, Ruikun, and Kaixin, for their friendship and great conversations. (Amrita, Sefah, Megha, and Albert, thanks for insisting that I should join you to watch the eclipse. It was simply spectacular.)

I am deeply grateful to be surrounded by such kind and genuine souls who have brought warmth over the years. Clare, thank you for being by my side, sharing life, laughter, and everything in between. Jinxing, Ruiqing, Qizhang, Xinyuan, Anke, Brian, Jiawei, and Jiaxin, thank you for bringing me countless moments of joy.

It marks the ninth year since I began my studies abroad. To my dearest family–Mom, Dad, and Aunt–thank you for being my rock and for your unwavering love and support, both emotional and financial. I am forever grateful to my grandfather, who never stopped expressing how proud he was of me. Though he is no longer here to witness this moment, his love remains, and *the spirit carries on*.

Dedication

To my parents.

"When your parents are alive, do not travel far; if you do, be sure to have a specific destination." – Confucius

Table of Contents

List of Figures xi

List of Tables xiii

1 Introduction 1

1.1 Epidemic models of measles 4

1.2 The vaccination game 6

1.3 Social structure, heterogeneity, and homophily 8

1.4 Early warning signals for critical transitions 9

1.5 Deep learning for time series classification 12

1.6 Thesis outline 14

2 Effect of homophily on coupled behavior-disease dynamics near a tipping point 16

2.1 Abstract 17

2.2 Introduction 17

2.3 Results 19

2.3.1 Model 19

2.3.2 Stochastic simulation 24

2.3.3 Homophily may lead to earlier disease outbreaks 25

2.3.4 Online events could significantly impact disease dynamics 28

2.3.5 Model explains the "early discussions" observed in social media data 30
2.4 Discussion . . . 33
2.5 Methods . . . 35
2.5.1 Measles outbreaks and cases . . . 35
2.5.2 Twitter data . . . 36
2.5.3 Sentiment analysis . . . 37

3 Forecasting infectious disease outbreak risks from vaccine sentiments on social media: A data-driven dynamical systems approach 39
3.1 Abstract . . . 40
3.2 Introduction . . . 40
3.3 Methods . . . 43
3.3.1 Deep learning model framework . . . 43
3.3.2 Theoretical data for training . . . 44
3.3.3 Empirical data for testing . . . 46
3.4 Results . . . 47
3.4.1 Performance of EWS and DL on theoretical data . . . 47
3.4.2 Performance of EWS and DL on empirical data . . . 49
3.5 Discussion . . . 53

4 Homophily in a multi-group vaccination game and its impact on disease dynamics 57
4.1 Abstract . . . 58
4.2 Introduction . . . 58
4.3 Methods . . . 60
4.4 Results . . . 65
4.4.1 Behavioral dynamics in the absence of diseases . . . 65
4.4.2 Impact of time-varying homophily levels under endemic conditions . 68
4.4.3 Time evolution of homophily-driven opinion polarization . . . 72
4.5 Discussion . . . 73

5 Conclusion **78**

5.1 Summary of findings 78

5.2 Future work 82

5.3 Concluding remarks 83

References **84**

APPENDICES **105**

A Supplementary Material: Effect of homophily on coupled behavior-disease dynamics near a tipping point **106**

A.1 Model derivation 106

A.2 Model simulation 109

A.3 Sensitivity analysis 110

B Supplementary Material: Homophily in a multi-group vaccination game and its impact on disease dynamics **113**

B.1 Derivation of the game dynamical equations 113

B.2 Matrix representation 115

B.3 Numerical methods 116

List of Figures

1.1 Measles cases and deaths in the United Kingdom from 1940 to 2023 4

1.2 Illustration of different social networks 9

1.3 Imitation processes in a population with two subgroups 10

2.1 Sequence of events in a directed graph 19

2.2 The ω–h parameter plane for the stability of the disease-free equilibrium with unanimous support for vaccination 25

2.3 Simulations of pro-vaccine sentiment percentage in the total population, infection prevalence, and susceptible percentage with different homophily levels 26

2.4 The first time when x drops below 95% and the first time when I increases above 10^{-4} in h–κ, δ–κ, and h–δ parameter planes 27

2.5 Possible realizations of the stochastic model with increasing magnitudes of noise or different proportions of social media users 29

2.6 Measles cases/infection percentage, number of Twitter posts, and selected EWS in three selected empirical datasets and one realization of the stochastic model 31

3.1 Two deep learning frameworks: CNN-LSTM and ResNet 44

3.2 Trends in two conventional EWS (variance and lag-1 autocorrelation) applied to simulated time-series data of x 48

3.3 Trends in two conventional EWS (variance and lag-1 autocorrelation) applied to simulated time-series data of corresponding T_p 49

3.4 ROC curves for predictions of an upcoming critical transition in theoretical time-series data 50
3.5 Trends in two conventional EWS and two deep learning classifiers applied to positive empirical Twitter data 51
3.6 Trends in two conventional EWS and two deep learning classifiers applied to neutral empirical Twitter data 52

4.1 Dynamics of pro-vaccine opinions and vaccine coverage under different levels of homophily in the general group when the disease is absent 67
4.2 Impact of homophily levels in the minority groups on final pro-vaccine proportion in the entire population when the disease is absent 68
4.3 Dynamics of pro-vaccine opinions across three groups under constant and logistically increasing homophily levels in the presence of the disease 70
4.4 Dynamics of pro-vaccine opinions across three groups under constant and logistically decreasing homophily levels in the presence of the disease . . . 71
4.5 Time evolution of Minority-Majority Divergence under varying basic reproduction numbers over the h_2–h_3 parameter plane 73
4.6 Pro-vaccine percentage in the entire population at $t = 200$ under varying basic reproduction numbers over the h_2–h_3 parameter plane 74

A.1 Simulations of the first time when the total pro-vaccine percentage x drops below the approximate herd immunity threshold for measles, and the starting time of the first epidemic 111
A.2 Corresponding elasticity coefficients of Figure A.1. 112

List of Tables

2.1 Parameters employed in Eqs. (2.6)–(2.11) and descriptions 23

2.2 Major measles outbreaks in the US, Canada and England, UK from 2014 to 2019 36

4.1 Parameters employed in Eqs. (4.1)–(4.5) and descriptions 64

Chapter 1

Introduction

Alongside warfare and technology, infectious diseases have played a significant role in shaping human history. In his Pulitzer-Prize-winning book *Guns, Germs, and Steel: The Fates of Human Societies*, Dr. Jared Diamond argues that Eurasia's historical dominance was due in part to the development of immunity to deadly infectious diseases after centuries of exposure to pathogens carried by domesticated animals, which was an advantage not shared by the native peoples of the New World [68]. Diseases such as smallpox served as "invisible weapons" and had a devastating impact during European colonization of the Americas and other regions [57, 58]. One striking example is the Spanish conquest of the Inca Empire. By the time Francisco Pizarro arrived in Peru in the early 16th century, smallpox had already decimated a large portion of the Inca population, including the emperor and his designated successor, igniting a civil war that left the empire politically fractured and militarily vulnerable [56].

Hundreds of years later, we no longer need to worry about suffering from smallpox, for it has been successfully eliminated through a global immunization program launched by the World Health Organization (WHO) [212, 114]. However, the threat of infectious diseases still remains [165]; apart from smallpox, no other infectious disease has been completely eradicated [82]. Measles, for example, continues to be a public health problem in many parts of the world, even though effective vaccines have been available for decades [206, 182]. More recently, the emergence of the COVID-19 pandemic pressed the pause button on the world, disrupting healthcare systems and economies, while also reshaping daily life for almost everyone–from the sudden shift to working from home to the widespread suspension of international travel.

The battle against infectious diseases has been a continuous thread throughout human

history, during which vaccines have emerged as a powerful tool for protection. They typically contain weakened or inactivated parts of a virus or bacteria, or generic instructions (such as mRNA) to prompt the body to produce a harmless piece of the pathogen. These components train the human immune system to recognize and respond quickly to real infections, thereby preventing illness [239, 45, 185]. In the past, when a highly contagious disease swept through a region, the population became infected with only two possible outcomes: recovery with acquired immunity, or death. These diseases spread quickly and then faded away just as rapidly, disappearing until another carrier of the pathogen arrived [176]. Vaccines offer a third path by stimulating the immune system to produce antibodies, which allow people to gain immunity without suffering from the disease.

Although vaccines have been proven effective in preventing diseases, motivations and concerns about vaccination differ between the public and individual levels. At the public level, high vaccination rates are essential to build herd immunity–when enough people become immune to a disease, either through previous infection or vaccination, the pathogen cannot spread freely in the population, protecting those who lack immunity [161]. At the individual level, the decision to get vaccinated is influenced by various factors, such as personal beliefs, fear of side effects, or access to vaccines. During the COVID-19 pandemic, many people expressed hesitation about the newly developed vaccines [74, 98, 225], worrying that the accelerated approval process might have overlooked potential risks. At the same time, proof of vaccination was often required for access to public places, which indirectly forced people to get vaccinated despite their concerns.

Individual decisions about vaccination are also affected by the behavior of others. People tend to rely on other people's actions to determine how they behave, especially in times of uncertainty. This is a common psychological and social phenomenon known as *social proof* [51]. When deciding whether to get vaccinated, people often look to those around them for guidance. If their family members, friends, or colleagues choose to get vaccinated, they will probably do the same since the vaccine gains credibility [188]. In contrast, if many in their social circle refuse to get vaccinated, they may be discouraged from doing so as well. In this sense, vaccination becomes a kind of "imitation game".

The spread of infectious diseases, the impact of vaccination, and imitation among individuals are closely interconnected and form a dynamic feedback loop. Specifically, people's tendency to imitate others' vaccination behavior can cause vaccine coverage in the population to either rise or fall, which, in turn, determines whether a pathogen persists or eventually dies out. Meanwhile, the state of the epidemic also affects people's incentive to get vaccinated: when a disease spreads quickly and leads to a large outbreak, the vaccination rate often surges in response; otherwise, people may feel less urgent and choose to delay or avoid vaccination [23]. This feedback loop unfolds within complex social networks,

where social connections vary in number and strength. Individuals with more social ties can influence, and be influenced by, more people, so their vaccination behavior is more likely to change. Moreover, clusters of like-minded individuals lead to more selective interactions, meaning people tend to encounter certain types of information while others are filtered out [15].

These intertwined dynamics also suggest that patterns in real-time behavioral data could potentially reflect underlying trends in vaccine uptake. Traditional surveillance systems typically rely on low-frequency and aggregate data, such as annual vaccination reports, which often suffer from reporting delays [222]. In recent years, the rise of digital platforms has enabled the collection of large-scale and high-frequency behavioral data. Social media activities, web searches, and other digital traces offer timely information on public sentiment on health-related topics, including attitudes about vaccination. Such data offer a promising way to improve the timeliness of vaccine coverage monitoring, allowing early detection of changes in vaccine uptake [192, 191].

This thesis addresses the following two main knowledge gaps:

(i) Individuals in a population are not identical, and there are differences in various aspects such as age and occupation, a phenomenon known as *heterogeneity*. These differences give rise to non-random contact patterns [116], where people are more likely to interact with others who share similar characteristics, a tendency referred to as *homophily* [157]. For example, children tend to interact more frequently with children in schools, while adults primarily socialize with adults in workplaces. Such contact patterns have been extensively studied in the context of disease transmission, but their influence on the spread of health-related behavior, such as vaccination, has received relatively little attention. In this thesis, we develop a generalized modeling framework that captures how individuals in different subgroups influence one another's vaccination behavior through imitation, and investigate how homophily affects both the behavior and disease dynamics.

(ii) Monitoring changes in vaccine coverage is of great importance. However, available vaccination rate data are often coarse in temporal resolution, making it difficult to capture timely feedback. This thesis aims to identify signals in high-frequency data that are potentially associated with trends in vaccine coverage, and to use deep learning (DL) models to predict whether the vaccination rate is approaching a tipping point. In situations when empirical data are insufficient to build a complete training dataset, we explore whether simulated data can serve as a substitute, a hybrid approach that combines mechanistic models and data-driven methods [14].

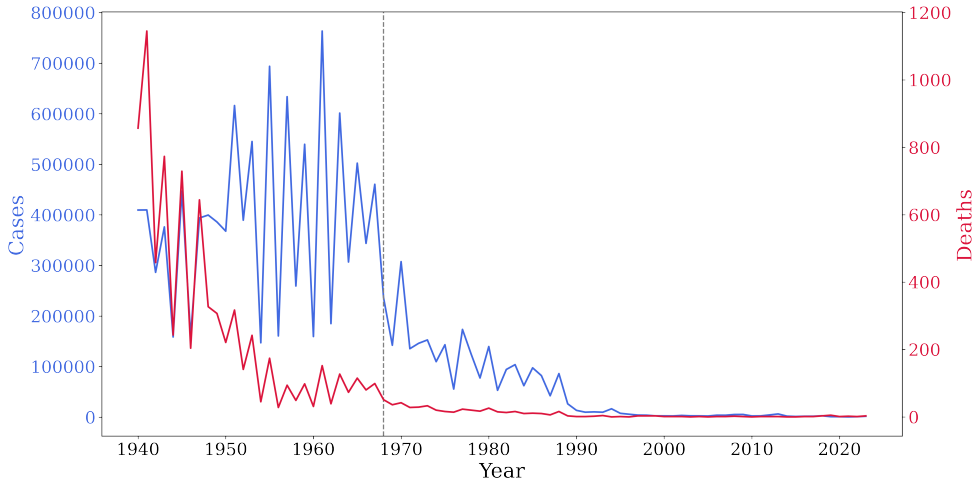


Figure 1.1: Measles cases (blue) and deaths (red) in the United Kingdom from 1940 to 2023. The gray dashed line marks 1968, the year the measles vaccine was introduced. Data from [220].

1.1 Epidemic models of measles

Among various infectious diseases, we chose measles as our focus based on several key factors. First, measles is one of the most contagious human diseases. This high transmissibility makes measles an ideal candidate for evaluating the accuracy and robustness of mathematical epidemic models [105]. Second, measles is highly preventable through vaccination. According to the Centers for Disease Control and Prevention (CDC), one dose of the measles, mumps, and rubella (MMR) vaccine is 93% effective against measles, and two doses increase the effectiveness to 97% [44]. However, measles outbreaks continue to occur occasionally due to decreases in vaccine coverage caused by vaccine hesitancy and behavioral factors, showing the critical role of human behavior in disease dynamics [183, 178]. Third, measles is supported by extensive historical epidemiological data, which facilitates model calibration and empirical validation. As shown in Figure 1.1, the United Kingdom has over 80 years of measles data. The sharp decline in cases and deaths after the introduction of measles vaccines in 1968 also highlights the strong impact of vaccines.

To study the transmission dynamics of measles, we adopt the Susceptible-Infectious-Recovered (SIR) model, one of the most fundamental compartmental frameworks in infectious disease modeling [131, 7, 130]. In its basic form, the SIR model divides the population into three compartments. Susceptible (S) individuals have no immunity to the disease and

may become infected upon exposure to the disease pathogen. Infectious (I) individuals are currently infected and capable of transmitting the disease pathogen to susceptible individuals. Recovered (R) individuals have acquired immunity and thus do not become infected again (if we assume lifelong immunity). Sometimes, the terms *infected* and *infectious* are interchangeable if we do not consider a latent period. In many cases, people who die from the disease are also included in the *recovered* class, as they no longer contribute to the transmission of the disease. This simplification is generally accepted when disease-related death is negligible with respect to the total population size.

We often use S, I, and R to denote the proportion of individuals in the corresponding compartments with $S + I + R = 1$. Their dynamics over time are described by a system of differential equations,

$$\frac{dS}{dt} = -\beta SI, \tag{1.1}$$

$$\frac{dI}{dt} = \beta SI - \gamma I, \tag{1.2}$$

$$\frac{dR}{dt} = \gamma I, \tag{1.3}$$

where β is the transmission rate and γ is the recovery rate (with $1/\gamma$ being the mean infectious period).

The system represented by Eqs. (1.1)–(1.3) captures only a single outbreak. This is because, in a closed population, the number of susceptible individuals decreases over time since they eventually become infected and recover, leaving too few for the pathogen to infect. Therefore, to model the resurgence of the disease, a continuous source of susceptible individuals must be introduced. Since measles primarily transmits among young children, and recovery from measles infection usually gives lifelong immunity, one common approach is to include natural births and deaths in the system, which gives

$$\frac{dS}{dt} = \Lambda - \beta SI - \mu S, \tag{1.4}$$

$$\frac{dI}{dt} = \beta SI - \gamma I - \mu I, \tag{1.5}$$

$$\frac{dR}{dt} = \gamma I - \mu R. \tag{1.6}$$

Here, Λ is the birth rate and μ is the natural death rate. For simplicity, it is often assumed that $\Lambda = \mu$ to maintain a constant population. We also introduce the concept of *basic reproduction number*, R_0, which represents the average number of secondary infections

caused by a single infectious individual in a fully susceptible population, given by $R_0 = \beta/(\gamma + \mu)$ [31]. For measles, the basic reproduction number is often cited between 12 and 18, which means that a person with measles can go on to infect 12 to 18 others, though the estimates could vary widely depending on the context [105].

While Eqs. (1.4)–(1.6) capture recurrent outbreaks of measles, it does not account for the seasonal variation often observed in real-world data. For childhood diseases like measles, the transmission rate tends to be higher during school terms due to increased contact among children [149]. This *seasonality* can be introduced by allowing the transmission rate β to vary periodically with time. For example, it can be modeled using a sinusoidal function $\beta(t) = \beta_0(1 + \beta_1 \sin(2\pi t))$, with β_0 the mean transmission rate and β_1 the amplitude of seasonal variation.

Although considering demographic and seasonal factors improves the practicability of the classic SIR model, these frameworks still assume that individuals behave passively in response to disease dynamics. In the post-vaccine era, modeling measles transmission requires incorporating behavioral dynamics, particularly vaccination decisions. Since measles vaccines are usually administered shortly after birth, it is reasonable to assume that vaccinated children enter the disease system as "recovered" individuals and unvaccinated children enter as susceptible individuals. If we let x denote the proportion of pro-vaccine parents in the population (since they are the ones who choose whether or not to vaccinate their children), μx represents the fraction of newborns who are vaccinated, while $\mu(1 - x)$ is the fraction that remains susceptible. Therefore, the system in Eqs. (1.4)–(1.6) becomes

$$\frac{dS}{dt} = \mu(1 - x) - \beta SI - \mu S, \tag{1.7}$$

$$\frac{dI}{dt} = \beta SI - \gamma I - \mu I, \tag{1.8}$$

$$\frac{dR}{dt} = \mu x + \gamma I - \mu R. \tag{1.9}$$

Naturally, x should be treated as a time-dependent variable, but how does it change as parents observe and imitate the vaccinating behavior of others?

1.2 The vaccination game

To model how vaccination behavior evolves, we turn to *game theory*, a mathematical framework used to analyze strategic decision-making where the outcome of a strategy depends on the actions made by all participants [97, 22]. In classical game theory, a *player* is an

individual who makes a decision, a *strategy* is the action a player chooses, and a *payoff* quantifies the resulting benefit or cost [174]. In the measles vaccination game, parents are the players who choose between two strategies: to vaccinate their child or not. The payoff depends on several factors, such as the perceived risk of the vaccine and the perceived risk of infection. Each parent aims to maximize their expected payoff by making their vaccination strategy. This choice determines the overall vaccine coverage in the population, which affects the infection probability for an unvaccinated child.

Classical game theory assumes that players are fully rational and make decisions with perfect information. However, in reality, parents cannot accurately assess the probability of infection. In addition, they tend to adopt new strategies by imitating others [20, 73], although their vaccination decision is locked in at the time of their child's birth. In response to these limitations, we adopt *evolutionary game theory*, which models how the distribution of strategies evolves within the population [155].

We follow the *replicator equation* that describes how the frequency of a strategy changes over time based on its *fitness* (analogous to *payoff* in classical game theory) relative to the average fitness in the population [118]. Here, the pro-vaccine strategy has a frequency of x and the anti-vaccine strategy has a frequency of $1-x$. We define $\pi_P(x)$ and $\pi_A(x)$ as the expected payoffs of pro-vaccine and anti-vaccine strategies, respectively, both of which are arbitrary functions of x. We let $\bar{\pi}(x) = x\pi_P(x) + (1-x)\pi_A(x)$ be the average payoff in the population. The replicator equation for the pro-vaccine strategy is

$$\frac{dx}{dt} = x[\pi_P(x) - \bar{\pi}(x)] = x(1-x)[\pi_P(x) - \pi_A(x)]. \tag{1.10}$$

An alternative way to derive the time evolution for x is to start with specific assumptions about the payoff functions [20]. We assume that the perceived risk of the vaccine is constant and the perceived risk of infection increases linearly with the current infectious percentage $I(t)$. Hence, the payoff difference would take the form $\Delta E = \pi_P - \pi_A = -c + mI$. We further assume that parents are willing to adopt a new strategy only when it offers a higher payoff, and their willingness to switch is higher if the payoff gain is larger. Given the sampling rate κ, anti-vaccine parents imitate pro-vaccine parents at a rate κx, and switch their strategy with a probability proportional to ΔE. The rate of change in the frequency of pro-vaccine parents becomes

$$\frac{dx}{dt} = \kappa x(1-x)\rho\Delta E \equiv \kappa x(1-x)(-\omega + I) \tag{1.11}$$

after rescaling, where ω is the relative vaccine risk. When $\Delta E < 0$, the payoff of the pro-vaccine strategy is lower, leading to a net shift to the anti-vaccine strategy in the

population. Eq. (1.11) remains valid as the sign of ΔE determines the direction of change. We shall notice that Eqs. (1.10) and (1.11) are equivalent since I is an implicit function of x through Eqs. (1.7)–(1.9).

A major limitation of the above assumptions is that when the infection level I is very low, the perceived risk of vaccination may exceed the risk of infection, making the anti-vaccine strategy more attractive regardless of x. As a result, x will decrease and eventually approach zero. However, vaccine refusal remains relatively rare for many childhood diseases, even in places with no infections. This suggests that other factors must be considered to help maintain a high level of vaccine support. One possible approach is to include the *injunctive social norm* that reflects socially encouraged behaviors [173],

$$\frac{dx}{dt} = \kappa x(1-x)[-\omega + I + \delta(2x-1)], \tag{1.12}$$

where δ quantifies the strength of the injunctive social norm. The term $\delta(2x-1)$ reinforces the majority strategy by being positive when $x > 0.5$ and negative when $x < 0.5$.

1.3 Social structure, heterogeneity, and homophily

As previously discussed in our research questions, real-world populations are not well-mixed, as individuals tend to interact through social networks shaped by spatial, cultural, and demographic factors. Here, we clarify how key terms are defined in this study. We refer to a population as *homogeneously mixed* if each individual has an equal probability of interacting with any other member. In contrast, a population is said to be *heterogeneously mixed* if there is a specific contact pattern due to factors such as geography, social structure, or individual characteristics. One important consequence of heterogeneous mixing is the phenomenon called *homophily*, the tendency of individuals to interact with similar others [157]. In other words, heterogeneity provides the basis for homophily, while homophily amplifies the effects of heterogeneity.

To illustrate contact patterns relevant to vaccination dynamics, we present three social networks in Figure 1.2, each consisting of 40 individuals with 60% supporting vaccination: (i) A homogeneously mixed network modeled as an Erdős-Rényi random graph [79], where individuals are randomly connected without contact patterns (Figure 1.2A). (ii) A two-block network, where each block contains 20 individuals who primarily interact within their own group. The pro-vaccine percentage is 60% in both groups, matching the entire population (Figure 1.2B). (iii) A polarized two-block network, where one group is dominated

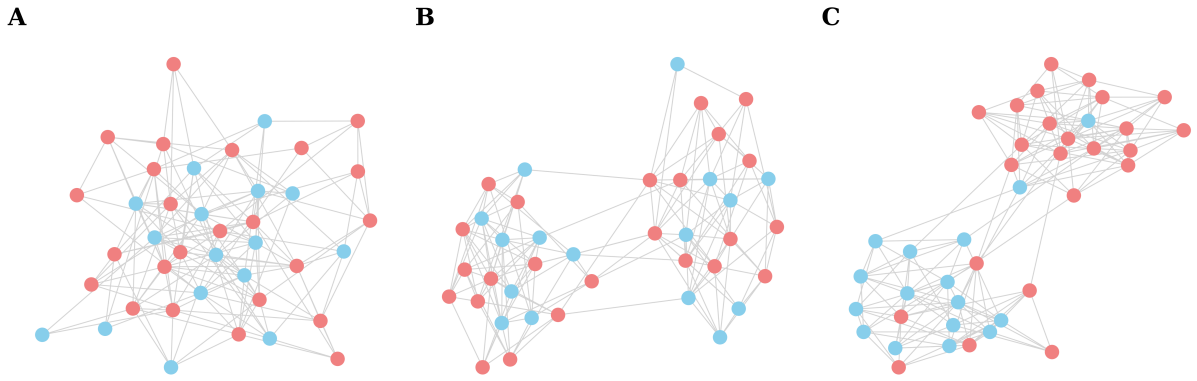


Figure 1.2: Illustration of different social networks. Red nodes represent pro-vaccine individuals, blue nodes represent anti-vaccine individuals, and edges denote social connections. (A) A homogeneously mixed network modeled as an Erdős-Rényi random graph. (B) A two-block structure, each containing a mix of pro- and anti-vaccine individuals in proportions matching the overall population. (C) A two-block structure with polarized vaccine sentiments.

by pro-vaccine individuals (90%) and the other is dominated by anti-vaccine individuals (30% pro-vaccine) (Figure 1.2C).

We now turn to the vaccination game. We do not focus on specific network topologies, as detailed behavioral data on person-to-person interactions are rarely available. Instead, we adopt an abstracted population structure by dividing the population into subgroups. We assume homogeneous mixing within each group, while the entire population exhibits heterogeneous mixing. In Figure 1.3, we consider a simplified case in which the population is divided into two subgroups. Within this framework, parents can imitate others within their own group or from the other group. Due to homophily, cross-group imitation occurs less frequently and has a weaker influence. We assume that parents are free to change their vaccine sentiment through imitation, but they do not switch between subgroups.

1.4 Early warning signals for critical transitions

Sudden and dramatic shifts are common features in many complex systems, ranging from the collapse of ecosystems [198, 87] and abrupt climate change [10] to crashes in the financial market [128]. These so-called *critical transitions* can occur when a system is gradually

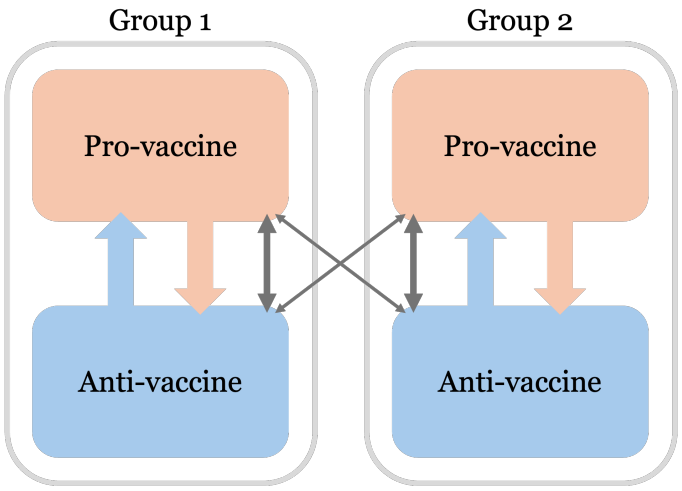


Figure 1.3: Imitation processes in a population with two subgroups. Bi-directional arrows represent imitation occurring both within and between groups, with arrow thickness indicating the relative strength of imitation.

pushed toward a *tipping point*, beyond which small perturbations trigger large and often irreversible changes. The popular book *The Tipping Point: How Little Things Can Make a Big Difference* by Malcolm Gladwell captures this idea in a broad societal context, showing how subtle behavioral changes can lead to widespread social transformations [100].

Such tipping behavior is particularly relevant in coupled human-environment systems, where interactions between natural dynamics and human decision-making create complex feedback loops [80]. In these systems, human behavior does not simply respond to environmental changes; it can actively alter key parameters of the system itself. For instance, policy decisions, social norms, or individual risk perceptions may influence the rate of pollution or resource consumption, which drives environmental changes [132]. A similar mechanism exists in the coupled behavior-disease model discussed in this study. When misinformation about vaccines spreads widely in a population, the perceived risk of vaccination can gradually increase. As a result, more individuals may opt out of vaccination because their subjective risk-benefit assessment has changed. These changes in collective behavior can lead to sharp declines in vaccination rates and, in extreme cases, to a collapse in vaccine uptake. A notable example is the MMR-autism scare in the late 1990s and early 2000s, triggered by a now-debunked study that falsely linked the MMR vaccine to autism [86].

The nonlinear nature of these coupled systems means that gradual changes in individual

behavior can lead to sudden shifts at the population level, such as from a disease-free state with full vaccine uptake to an endemic state with zero vaccine uptake. Mathematically, this kind of critical transition can be understood through the lens of *bifurcation theory* in dynamical systems, which describes how small changes in parameters may induce qualitative shifts in the long-term behavior of the system [196, 136]. *Bifurcations* are generally classified as local and global, depending on whether the qualitative change can be detected by a stability analysis of the equilibria. This study focuses on *local bifurcations*, which occur when the stability of an equilibrium changes as eigenvalues of the system's Jacobian matrix cross the imaginary axis [213, 138]. For the coupled behavior-disease system described in Eqs. (1.7)–(1.9) and Eq. (1.12), the point $(S^*, I^*, R^*, x^*) = (0, 0, 1, 1)$ represents a disease-free equilibrium with full vaccine uptake, which is stable when the social norm is stronger than the relative vaccine risk ($\delta > \omega$) [173]. However, as ω increases and exceeds δ due to the spread of misinformation, this equilibrium destabilizes, and the system converges to a new steady state: an endemic equilibrium with zero vaccine uptake. The threshold $\omega = \delta$ marks a critical transition, beyond which a collapse of vaccine uptake occurs [175].

Importantly, the approach to a bifurcation is often accompanied by changes in the system's properties. One widely studied phenomenon is called *critical slowing down* (CSD), which refers to the system's increasing recovery time from small perturbations as it approaches a tipping point [235]. Imagine a ball placed at the bottom of a well. If the well is deep and steep, a small nudge will cause the ball to quickly roll back to the bottom, just as a system returns to equilibrium after a small perturbation. However, as the system approaches a tipping point, the well becomes shallower and flatter, and the same nudge causes the ball to return more slowly [62]. This slower recovery reflects the reduced resilience of the system.

Although CSD can theoretically be detected by applying small disturbances to a system and measuring how quickly it returns to its initial state, this method is impractical in real-world settings. Moreover, CSD occurs because the dominant eigenvalue of the system approaches zero from below [221], but the governing equations of a real-world dynamical system are usually unknown. As a result, researchers have turned to statistical techniques that can capture the signs of CSD in time series data. Two commonly used indicators are lag-1 autocorrelation and variance [197].

As a system approaches a tipping point, its recovery from disturbances becomes slower, causing its current state to resemble its previous state more closely. Lag-1 autocorrelation measures the correlation between the state of the system at time t and time $t-1$, and tends to increase toward 1 when a system experiences CSD. Meanwhile, reduced resilience makes the system more sensitive to random fluctuations, allowing them to push the system

further from equilibrium and leading to an increase in variance. Therefore, increases in lag-1 autocorrelation and variance are considered signals of CSD, which we often refer to as *early warning signals* (EWS). To detect such signals, the trend of an indicator is typically quantified using Kendall's tau correlation coefficient, which ranges from -1 (strictly decreasing) to 1 (strictly increasing) [64].

In addition to lag-1 autocorrelation and variance, other indicators of CSD have also been proposed in recent studies, such as spectral reddening and detrended fluctuation analysis [62]. While these statistical indicators and trend detection methods have proven useful in many settings, they often struggle with noisy or incomplete data and rely on strong assumptions about system dynamics. These limitations have motivated a growing interest in data-driven approaches for detecting EWS in time series data. In the next section, we explore how deep learning models can be used for time series classification in the context of critical transitions.

1.5 Deep learning for time series classification

From IBM's Deep Blue defeating a world chess champion in the 1990s to DeepMind's AlphaGo mastering the game of Go in 2016, artificial intelligence (AI) has made remarkable progress, not just in game-playing, but in how machines learn and make decisions. AI broadly refers to the development of algorithms that enable machines to perform tasks that typically require human intelligence. Within this field, machine learning (ML) has emerged as a dominant paradigm that applies data-driven methods to automatically learn patterns and make predictions without being explicitly programmed [124]. A key subset of ML is DL, which employs multi-layered neural networks to reveal more complex relationships in data [140, 103]. Neural networks are computational models inspired by the structure of the human brain, composed of layers of interconnected units ("neurons") that learn to extract features from input data.

While DL has achieved significant success in image recognition and natural language processing, its application to classical time series analysis is a relatively new but rapidly growing area of research. Researchers have demonstrated that deep neural networks can effectively capture patterns from sequential data, making them suitable for time series classification tasks [123]. In time series classification, the goal is to assign a label to an entire sequence based on its temporal characteristics. This is distinct from time series forecasting, which aims to predict future values based on past observations. In the context of EWS detection, classification allows models to recognize whether a system is in a stable regime or approaching a critical transition, without explicitly modeling future trajectories.

Among various deep learning architectures, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) are particularly effective in dealing with sequential data. CNNs were originally developed for visual data [141, 135], but have since been adapted for one-dimensional time series by treating the temporal axis analogously to the spatial axes in images [133]. The core idea behind CNNs is the use of convolutional layers that apply learnable filters across segments of the input. As these filters move along the sequence, they generate *feature maps* that capture important temporal features. Convolutional layers are typically followed by pooling operations that reduce the size of the feature maps while preserving key information, and then by fully-connected layers that combine learned features for final classification.

RNNs are designed to model temporal sequences by maintaining a form of "memory" throughout time steps, enabling them to capture dependencies within sequential data. This is achieved through recurrent connections, where the output at each time step is fed back into the network as input for the next step. However, standard RNNs face difficulties in learning long-range dependencies due to exploding or vanishing gradients [26]. To combat this problem, Long Short-Term Memory (LSTM) networks are introduced, which feature a separate memory state that persists across time steps without being repeatedly transformed by activation functions or weighted connections [117]. Detailed architectures and implementation examples of 1D-CNN and LSTM can be found in [201].

Standard CNNs also suffer from exploding or vanishing gradients as the number of layers increases. A Residual Neural Network (ResNet) addresses this issue by incorporating skip connections that directly link the output of one layer to a deeper layer, bypassing the layers in between. Each skip connection, along with the intermediate layers it bypasses, forms a *residual block*. A ResNet is constructed by stacking multiple such residual blocks, allowing the training of very deep networks without performance degradation [108]. ResNet has demonstrated success in image recognition and has more recently been adapted for time series classification tasks [230].

Building on these architectures, researchers have begun to explore the potential of deep learning in the detection of EWS before critical transitions [38, 66, 37]. Since critical transitions are often related to local bifurcations, and the center manifold theorem suggests that high-dimensional systems can be reduced to low dimensions near these bifurcation points, it becomes feasible to focus on simplified normal forms of bifurcations. By simulating these systems, one can generate time series data that capture both CSD before critical transitions and stable dynamics where no transition occurs, which allows the creation of labeled datasets where the timing and type of bifurcation are known. DL models trained on such data can learn to identify early patterns specific to different bifurcation types, while also distinguishing them from baseline scenarios where no transition is going to happen.

Notably, [38] shows that the CNN-LSTM framework can provide EWS in empirical and model time series in ecology and epidemiology, with greater sensitivity and specificity than traditional statistical methods.

1.6 Thesis outline

The remainder of this thesis is organized as follows. In Chapter 2, we introduce a novel coupled behavior-disease model that distinguishes between social media users and non-users when modeling imitation dynamics on vaccination behavior. Unlike existing coupled models that assume imitation dynamics occur in a homogeneous population, our model explicitly incorporates heterogeneity in social media engagement and includes homophily to reflect the tendency of individuals to interact within their own social groups. Our analysis shows that increasing homophily prevents the spread of pro-vaccine strategies, making the population more vulnerable to crossing a tipping point into an endemic state with no vaccine uptake. Furthermore, EWS such as increased variance and lag-1 autocorrelation were observed in social media data before the 2014 measles outbreak in California [175]. By analyzing additional datasets, we find this to be a common phenomenon. Our mathematical model reproduces similar early earning patterns in simulated social media data, reflecting the mechanisms that may be at the heart of these empirical observations. This chapter has been published in *Mathematical Biosciences* [109].

In Chapter 3, we present a data-driven approach for detecting EWS of vaccine sentiment shifts that may precede disease outbreaks. Instead of relying on low-resolution vaccination rate data, our method analyzes real-time changes in vaccine sentiment derived from high-frequency social media data to forecast outbreak risk. In addition to the widely employed CNN-LSTM architecture in previous studies on DL-based EWS [38, 37], we also develop a ResNet classifier to compare performance. Both classifiers are trained on time series of social media activity generated from the stochastic coupled behavior–disease model introduced in Chapter 2. Lévy noise, rather than Gaussian white noise, is used in order to simulate heavy-tailed fluctuations commonly observed in real-world social and epidemiological systems. We compare the performance of DL classifiers with traditional statistical EWS based on the theory of CSD, such as variance and lag-1 autocorrelation, on both theoretical and empirical disease-related social media data. Our work demonstrates the potential of using deep learning with real-time social media data to forecast emerging public health risks with greater accuracy and interpretability. It also provides indirect evidence supporting the feasibility of training DL models with simulated data when real-world data is highly limited. This chapter is under review at *Mathematical Biosciences*

and Engineering.

In Chapter 4, we extend the investigation in Chapter 2 by modeling homophily in a more generalized multi-group vaccination game. In contrast to binary grouping, this model divides the population into an arbitrary n distinct groups, allowing for more grouping possibilities and considering more complex cross-group interactions. We narrow our focus to a three-group game, where we divide the population into one majority group representing the general public and two minority groups: a strongly pro-vaccine group and a strongly anti-vaccine group. The study reveals that homophily can lead to multi-strategy equilibria within each group, causing opinion polarization at the population level, even when the disease is highly transmissible. Our model provides a possible explanation for a real-world observation: even in regions with high overall vaccination coverage, there can exist clusters of communities with persistently low vaccine uptake. This phenomenon cannot be captured by previous models that assume homogeneous mixing. To our knowledge, this is the first study to rigorously analyze the equilibrium structure in a multi-group vaccination game under homophily. This chapter has been submitted to *Proceedings of the Royal Society B: Biological Sciences*.

Finally, in Chapter 5, we conclude the thesis by synthesizing the main findings and discussing their broader implications for behavior-disease modeling and public health practice. We also discuss the key limitations of our work and propose several directions for future research.

Chapter 2

Effect of homophily on coupled behavior-disease dynamics near a tipping point

This chapter is based on the paper: Z. He and C. T. Bauch. Effect of homophily on coupled behavior-disease dynamics near a tipping point. *Mathematical Biosciences*, 2024, 376:109264.

2.1 Abstract

Understanding the interplay between social activities and disease dynamics is crucial for effective public health interventions. Recent studies using coupled behavior-disease models assumed homogeneous populations. However, heterogeneity in population, such as different social groups, cannot be ignored. In this study, we divided the population into social media users and non-users, and investigated the impact of homophily (the tendency for individuals to associate with others similar to themselves) and online events on disease dynamics. Our results reveal that homophily hinders the adoption of vaccinating strategies, hastening the approach to a tipping point after which the population converges to an endemic equilibrium with no vaccine uptake. Furthermore, we find that online events can significantly influence disease dynamics, with early discussions on social media platforms serving as an early warning signal of potential disease outbreaks. Our model provides insights into the mechanisms underlying these phenomena and underscores the importance of considering homophily in disease modeling and public health strategies.

2.2 Introduction

Measles, once a highly contagious disease that claimed millions of lives annually, underwent a significant decline in incidence following the development of a vaccine in the 1960s [77, 237]. This decline was further accelerated by the introduction of highly effective immunization programs [237], especially after the introduction of a second vaccine dose in the 1990s [183]. In the year 2000, the United States proudly declared measles elimination, having experienced no transmission for a minimum of 12 months [43, 179]. Since then, over 20 million deaths have been prevented globally through measles vaccination [61]. However, the elimination of measles in the US proved to be short-lived, and worldwide elimination remains elusive. The constant inflow of the virus carried by international travelers from measles-endemic regions [142, 177], coupled with vaccine refusal leading to reduced immunization rates, creates ideal conditions for resurgence to occur.

One of the primary factors contributing to measles vaccine refusal is the concern regarding vaccine safety. In 1998, a series of articles were published in the esteemed medical journal *The Lancet*, falsely linking the measles-mumps-rubella (MMR) vaccine to autism among children [226]. Despite being discredited as fraudulent [195], these articles had a lasting negative impact on measles vaccination rates worldwide [172, 25]. During the early 2000s, there was a significant decrease in public trust regarding the safety of the MMR vaccine, resulting in a notable decline in the vaccination rate in England between 2000

and 2004 [216]. The measles outbreak that occurred in Disneyland, California, from December 2014 to February 2015 also highlighted the consequences of vaccine refusal and its association with disease outbreaks [150, 183]. The California Department of Public Health reported a substantial decrease in MMR vaccine coverage among kindergarten students in California from 2010 to 2014 [40].

Human behaviors both influence and are influenced by disease dynamics, as outbreaks often trigger behavioral responses that effectively prevent the spread of the disease [91, 90]. To explore this interrelationship, it is essential to consider human behaviors in epidemic models. Evolutionary game theory provides a theoretical framework for understanding the dynamic nature of individual behaviors. Since the pioneering work by Fine and Clarkson [85], there has been a significant interest in applying evolutionary game theory to model this two-way feedback in adopting vaccinating strategies [93, 223, 22, 20, 193, 28], with some research focusing particularly on measles [21, 19, 72, 205, 173]. Later progress included considering more than one subgroup of individuals [189] and social contact networks [169]. Recent studies have expanded upon this approach by incorporating bifurcation theory and placing considerable emphasis on the identification of early warning signals (EWS) for critical transitions [210, 197, 63, 126]. When a system approaches a tipping point, even small perturbations can trigger a shift to an alternative state, resulting in a reduced ability to recover from perturbations. This characteristic behavior is referred to as critical slowing down (CSD). Certain EWS, such as increased variance and lag-1 autocorrelation, serve as reliable indicators for disease emergence [32]. They enable researchers to proactively alter the future trajectory of a system, even when working with imperfect epidemiological data [33].

Monitoring the vaccination rate trend serves as an effective means of evaluating the behavior-disease model in real-world scenarios. However, a notable challenge arises from the limited availability of vaccination data that was typically obtained annually. For instance, National Immunization Survey (NIS) monitors vaccination coverage among children through annual phone surveys [42]. The coarse-scale nature of the time series data poses difficulties for researchers in accurately tracking trends, identifying the specific timeframe when vaccine coverage falls below critical thresholds, and predicting potential future outbreaks.

The widespread use of social media platforms has significantly facilitated the dissemination of both pro- and anti-vaccine sentiments. Notably, influential figures on platforms like Twitter have had a profound impact on shaping public opinion regarding vaccination [25]. Moreover, the availability of high-frequency social media data has introduced novel avenues for studying the dynamics of infectious diseases. For instance, Pananos *et al.* [175] conducted research that involved geocoding Tweets related to the MMR vaccine prior to

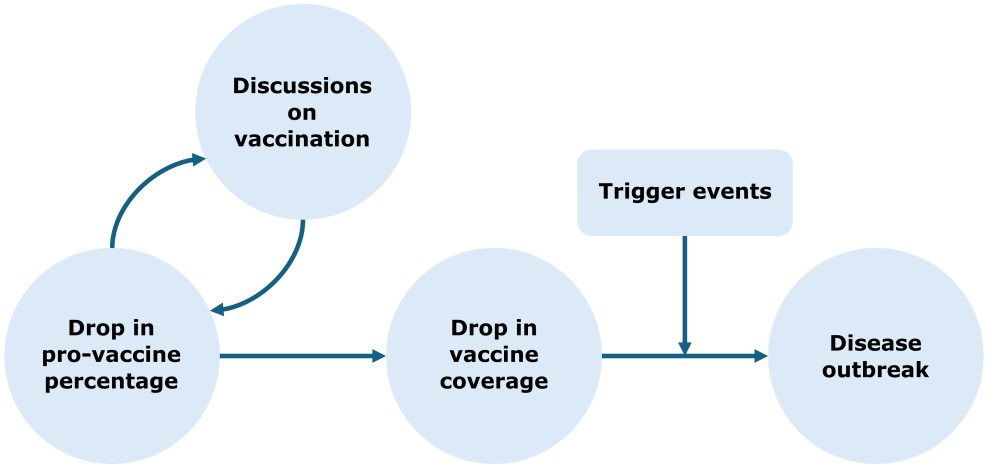


Figure 2.1: Sequence of events in a directed graph.

the measles outbreak in Disneyland, California from 2014 to 2015. By observing EWS derived from the time-series data of these Tweets, the study reveals CSD in the dynamics of vaccine-related discussions.

This article presents an extension of the coupled behavior-disease model, incorporating two different social groups, namely social media users and non-users. Our model captures the imitation processes occurring within each group and between different groups [111]. The underlying causal inference is illustrated in Figure 2.1. We hypothesize that a decrease in the proportion of pro-vaccine individuals will lead to a decline in vaccine coverage, thereby creating the conditions necessary for a potential outbreak. Furthermore, a decrease in the pro-vaccine percentage (in a population originally dominated by pro-vaccine opinion) is assumed to foster a more balanced distribution of opinions, stimulating discussions on vaccination. The coupled dynamics allow us to estimate online discussions with a relatively straightforward formula, which can be validated qualitatively against real-world social media data.

2.3 Results

2.3.1 Model

To address our objectives, we propose a behavior-disease model that couples disease transmission and vaccinating behavior dynamics. Our approach commences with the widely used

Susceptible-Infected-Recovered (SIR) epidemic model that incorporates natural births and deaths [7], where individuals are classified as susceptible, infected, or recovered based on their health conditions during an epidemic. To capture homophily, we categorize each individual as either a social media user or a non-user and assume that their social group will not change; to capture the dynamics of vaccinating behaviors, we employ evolutionary game theory [20, 21, 173], and categorize each individual as either pro-vaccine or anti-vaccine, with the freedom to observe and adopt the vaccination strategies of others, subsequently adjusting their own strategies. However, due to homophily, interactions between individuals from different social groups occur less frequently [157], and thus are less influential in shaping opinions. We assume that individual behaviors can be spread and adopted quickly [46], and that individuals have a 100% commitment to their vaccination strategies. Consequently, people's attitudes towards vaccination evolve on a similar timescale as interepidemic periods [232, 207]. At each time point t, the number of new individuals adopting vaccination (vaccinators) entering the system is proportional to the total number of pro-vaccine individuals present at that specific time. Likewise, the number of new non-vaccinators entering the system follows the same relationship. Furthermore, we assume lifetime immunity after vaccination, so vaccinators enter the system as recovered individuals, while non-vaccinators enter as susceptible individuals.

We let x_1 (respectively, x_2) be the proportion of frequent social media users (respectively, social media non-users, or infrequent users) that are pro-vaccine. Hence, the proportion of anti-vaccine individuals among frequent social media users (respectively, non-users or infrequent users) is $1-x_1$ (respectively, $1-x_2$). We assign the payoffs $E_i(P)$ and $E_i(A)$ to pro-vaccine and anti-vaccine individuals, respectively, within social group i as [20, 173]

$$E_i(P) = -r_v + \delta_0 f_i^P(x_1, x_2), \tag{2.1}$$

$$E_i(A) = -r_{\mathrm{n}}(I(t)) + \delta_0 f_i^A(x_1, x_2), \tag{2.2}$$

where $E_i(S)$ denotes the payoff of adopting the strategy S in the social group i. r_v is the perceived probability of significant morbidity from the vaccine, while r_{n} is the perceived probability of significant morbidity from being infected, which is assumed to be a function of disease prevalence $I(t)$ at time t. δ_0 quantifies the effect of injunctive social norms (behaviors that individuals are expected to follow and expect others to follow, representing what people should do, even if it goes against their immediate interests) [52], f_i^S is a function representing the group pressure of adopting the strategy S in the social group i [112, 113].

We make the assumption that individuals will only imitate others if switching their vaccination strategy leads to a higher payoff [20]. Moreover, the imitation rate is assumed

to be proportional to the payoff gain. Therefore, when $E_i(P) > E_i(A)$, the rate at which anti-vaccine individuals in the social group i imitate pro-vaccine individuals within the same social group can be determined by calculating

$$\kappa x_i(1-x_i)[E_i(P)-E_i(A)], \tag{2.3}$$

in which κ is the sampling rate, and represents the rate of social learning. When $E_i(A) > E_i(P)$, $E_i(P) - E_i(A)$ becomes negative, making Eq. (2.3) essentially

$$-\kappa(1-x_i)x_i[E_i(A)-E_i(P)]. \tag{2.4}$$

The negative sign in the front denotes the population outflow, which indicates that anti-vaccine individuals remain steadfast in their strategy, while pro-vaccine individuals are switching to an anti-vaccine strategy. A similar calculation can be applied to assess the imitation dynamics between individuals from different social groups. However, both the imitation rate and the social norm strength between social media users and non-users are attenuated by a decay factor of $(1-h)$, where h ranges between 0 and 1, representing the level of homophily [157, 159]. A value of $h = 1$ indicates complete homophily, wherein individuals exclusively interact with similar others, so the social norm from the other group will have no impact. Conversely, when $h = 0$, individuals have no preference for whom they interact with, so they imitate others at the same rate and are equally influenced by the social norm in both groups. We obtain the time evolution of x_i:

$$\begin{aligned}\frac{dx_i}{dt} &= \kappa x_i(1-x_i)[E_i(P)-E_i(A)] \\ &\quad + (1-h)\kappa x_j(1-x_i)\max[E_j(P)-E_i(A),0] \\ &\quad - (1-h)\kappa(1-x_j)x_i\max[E_j(A)-E_i(P),0],\end{aligned} \tag{2.5}$$

where $i, j \in \{1, 2\}$ and $i \neq j$ since we only consider two social groups. The group pressure functions in Eqs. (2.1) and (2.2) are

$$\begin{aligned}f_i^P &= x_i + (1-h)x_j, \\ f_i^A &= 1 - x_i + (1-h)(1-x_j),\end{aligned}$$

with which Eq. (2.5) becomes

$$\begin{aligned}\frac{dx_i}{dt} &= \kappa x_i(1-x_i)[-r_v + cI + \delta_0(2x_i - 1 + (1-h)(2x_j-1))] \\ &\quad + (1-h)\kappa\Big\{x_j(1-x_i)\max[-r_v + cI + \delta_0(2-h)(x_i+x_j-1),0] \\ &\quad - (1-x_j)x_i\max[r_v - cI - \delta_0(2-h)(x_i+x_j-1),0]\Big\}.\end{aligned}$$

After substituting $\omega \equiv r_v/c$, $\delta \equiv \delta_0/c$, and $\kappa \equiv \kappa c$, we obtain the coupled model defined by the following system of differential equations:

$$\frac{dS}{dt} = \mu(1-x) - \beta SI - \mu S, \tag{2.6}$$

$$\frac{dI}{dt} = \beta SI - \gamma I - \mu I, \tag{2.7}$$

$$\frac{dR}{dt} = \mu x + \gamma I - \mu R, \tag{2.8}$$

$$\begin{aligned}\frac{dx_1}{dt} &= \kappa x_1(1-x_1)[-\omega + I + \delta(2x_1 - 1 + (1-h)(2x_2-1))] \\ &\quad + (1-h)\kappa\Big\{x_2(1-x_1)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\ &\quad - (1-x_2)x_1\max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\},\end{aligned} \tag{2.9}$$

$$\begin{aligned}\frac{dx_2}{dt} &= \kappa x_2(1-x_2)[-\omega + I + \delta((1-h)(2x_1-1) + 2x_2 - 1)] \\ &\quad + (1-h)\kappa\Big\{x_1(1-x_2)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\ &\quad - (1-x_1)x_2\max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\},\end{aligned} \tag{2.10}$$

where S, I, and R are the proportions of susceptible, infectious, and recovered individuals respectively. The proportion of pro-vaccine individuals in the entire population, x, is calculated as

$$x = \frac{x_1 N_1 + x_2 N_2}{N_1 + N_2}. \tag{2.11}$$

N_1 and N_2 are assumed to be constant over time. The differential equation for R can be eliminated as $S + I + R = 1$ always holds. The list of parameters employed in Eqs. (2.6)–(2.11) is presented in Table 2.1. For a detailed elaboration of our model, please refer to Appendix A.1.

We assume that the rate of pro-vaccine sentiments expressed on social media posts closely aligns with the actual current vaccination rates in infants [194], allowing the number of social media posts to serve as a measurement of vaccination behaviors within the system. Here, we further devise a formula that captures the following characteristics for calculating the number of social media posts:

(i) When the disease is not transmitted in the population ($I \approx 0$, very small), a substantial number of social media posts (regardless of sentiment) should be generated when

Parameter	Description	Baseline value	References
N_1	Pro-vaccine frequent social media users	1000	Choice of interest [181, 70]
N_2	Pro-vaccine infrequent/non- social media users	19000	Choice of interest [181, 70]
R_0	Basic reproduction number	16	[7, 105]
γ	Recovery rate	$365/14$ yr^{-1}	[7, 77]
μ	Natural birth (and death) rate	$1/60$ yr^{-1}	[7, 77]
β	Transmission rate	417 yr^{-1}	[7, 77]
κ	Sampling rate	3000 yr^{-1}	[21, 173, 175, 38]
ω	Relative vaccine risk	2×10^{-4}	[21, 173, 175, 38]
δ	Social norm strength	3×10^{-4}	[21, 173, 175, 38]
h	Homophily level	0.8	Choice of interest

Table 2.1: Parameters and descriptions. Population sizes N_1 and N_2 are hypothetical but the ratio of social media users is obtained from [181, 70]. N_1 includes high activity users only while N_2 also includes non-users. The transmission rate β is calculated by $\beta = R_0(\gamma + \mu)$. The relative vaccine risk is assumed to increase linearly from the baseline value $\omega = (2 + 0.5t) \times 10^{-4}$.

polarization emerges (i.e. the proportions of individuals holding different opinions are relatively balanced, with x_1 and $1 - x_1$ being approximately equal). If one opinion overwhelmingly dominates, a prolonged debate becomes unnecessary [121].

(ii) When I increases, a greater number of posts (regardless of sentiment) should be observed online. This aligns with the notion that online discussions on vaccination tend to surge during disease outbreaks [175].

One possible way to calculate the number of pro-vaccine social media posts is

$$T_p = cx_1(1 - x_1) + \tau I, \tag{2.12}$$

where c and τ are scaling factors that characterize the impact of opinion difference and disease prevalence, respectively. T_p reaches its maximum value when $x_1 = (1 - x_1) = 0.5$, indicating a perfectly polarized population [164], particularly when I is small. Similarly, the number of anti-vaccine posts can be calculated as

$$T_a = c'x_1(1 - x_1) + \tau' I \tag{2.13}$$

with perhaps different scaling factors c' and τ'. However, as will be discussed in the subsequent sections, the number of collected anti-vaccine posts is significantly lower than that of pro-vaccine posts. This discrepancy leads to a less precise estimation. Consequently, our project will concentrate solely on analyzing the number of pro-vaccine posts.

2.3.2 Stochastic simulation

The deterministic model shown in Eqs. (2.6)–(2.10) is converted into a stochastic model driven by an α-stable Lévy process,

$$d\mathbf{X} = \mathbf{f}(\mathbf{X}, t)dt + \boldsymbol{\sigma} dL, \tag{2.14}$$

where $\mathbf{X} = [S, I, x_1, x_2]$, $\mathbf{f}(\mathbf{X}, t)$ is a column vector of the right-hand sides of Eqs. (2.6), (2.7), (2.9), and (2.10). dL is an additive α-stable Lévy process, and $\boldsymbol{\sigma} = [\sigma_S, \sigma_I, \sigma_{x_1}, \sigma_{x_2}]$ is a column vector of scaling factors controlling the magnitude of noise for S, I, x_1, and x_2.

It is noteworthy to highlight our use of an α-stable Lévy process in generating additive noises, rather than a conventional Wiener process as in many studies on stochastic models. The purpose is to better estimate rare and extreme events, such as a drastic change in the pro-vaccine percentage due to influential news reports or social media posts by celebrities. These events often have an infinitesimal probability of occurrence, but would hugely impact the dynamics of the system, and thus predicting them with high precision is of great importance [158]. In reality, such events might occur at a higher frequency than a normal random variable predicts. As a result, a random variable under a fat-tailed distribution is desired [170]. We choose an α-stable Lévy process with stability parameter $\alpha \in (0, 2)$ because its tail decays like a power law. Furthermore, an α-stable Lévy process possesses independent and stationary increments, which provide convenience in designing algorithms for simulation.

We introduce a modification to the formula governing the count of pro-vaccine social media posts by adding a noise term with

$$T = mT_p + k\xi. \tag{2.15}$$

Here, m and k serve as scaling factors, while ξ represents an additive noise that follows a skewed fat-tailed stable distribution, $\xi \sim \text{Stable}(\alpha, \beta, c = c(T_p), \mu = \mu(T_p))$ with stability parameter $\alpha \in (1, 2)$ and skewness parameter $\beta \in (0, 1)$. The scale parameter c and the location parameter μ should depend on T_p. Specifically, both c and μ should increase as T_p increases (see Appendix A.2). The underlying rationale is that a surge in T_p signifies an elevated prominence of vaccination discussions, thereby elevating the likelihood of re-posting and, consequently, substantially augmenting the total post count.

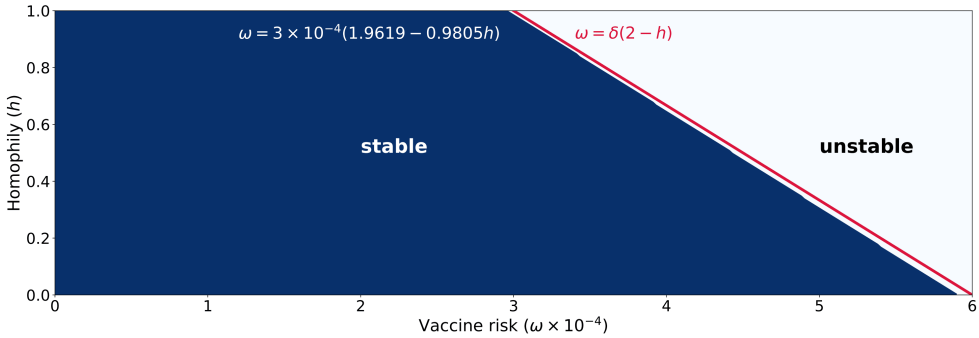


Figure 2.2: The ω–h parameter plane for the stability of the disease-free equilibrium ($\mathcal{E}$) with unanimous support for vaccination $(S, I, x_1, x_2) = (0, 0, 1, 1)$. $\mathcal{E}$ is stable in the dark blue region on the left and unstable in the light blue region on the right. Red line is the estimated critical transition $\omega = \delta(2 - h)$.

2.3.3 Homophily may lead to earlier disease outbreaks

Our investigation reveals that homophily may lead to earlier disease outbreaks by preventing the spread of pro-vaccine strategies across different social groups. The disease-free equilibrium with unanimous support for vaccination (leading to full vaccine uptake) $\mathcal{E} = (S, I, x_1, x_2) = (0, 0, 1, 1)$ is of particular interest as this state represents the desired outcome. We generated plots of the stability of such an equilibrium when we varied the perceived vaccine risk ω, while considering different levels of homophily h. Prior research has discussed that a critical transition from full to zero vaccine uptake happens when the perceived vaccine risk ω exceeds the social norm strength δ if different subgroups are not considered in a population, which is identical to complete homophily ($h = 1$) [173, 175]. Due to the complexity of Eqs. (2.9) and (2.10), we were unable to derive an analytic expression describing the relationship between the critical transition point and the level of homophily, but after a number of experiments, we expected it to be $\omega = \delta(2 - h)$, which is the time when $-\omega + I + \delta(2 - h)(x_1 + x_2 - 1) = 0$ at $I = 0$ and $x_1 = x_2 = 1$.

In our simulation where we maintained a fixed social norm strength $\delta = 3 \times 10^{-4}$, and set initial conditions $(S(0), I(0), x_1(0), x_2(0)) = (0.01, 10^{-6}, 0.99, 0.99)$, we observed that such a critical transition occurred when the perceived vaccine risk ω was slightly smaller than the social norm strength. In fact, the regression line is $\omega = (-2.9414h + 5.8858) \times 10^{-4}$, which is close enough to $\omega = (-3h + 6) \times 10^{-4}$ as we expect. The difference occurs because of the small magnitude of I and the fact that $x_1 + x_2 - 1$ is slightly less than 1. From Figure 2.2, we conclude that $\omega > \delta(2 - h)$ is a sufficient but not necessary condition for

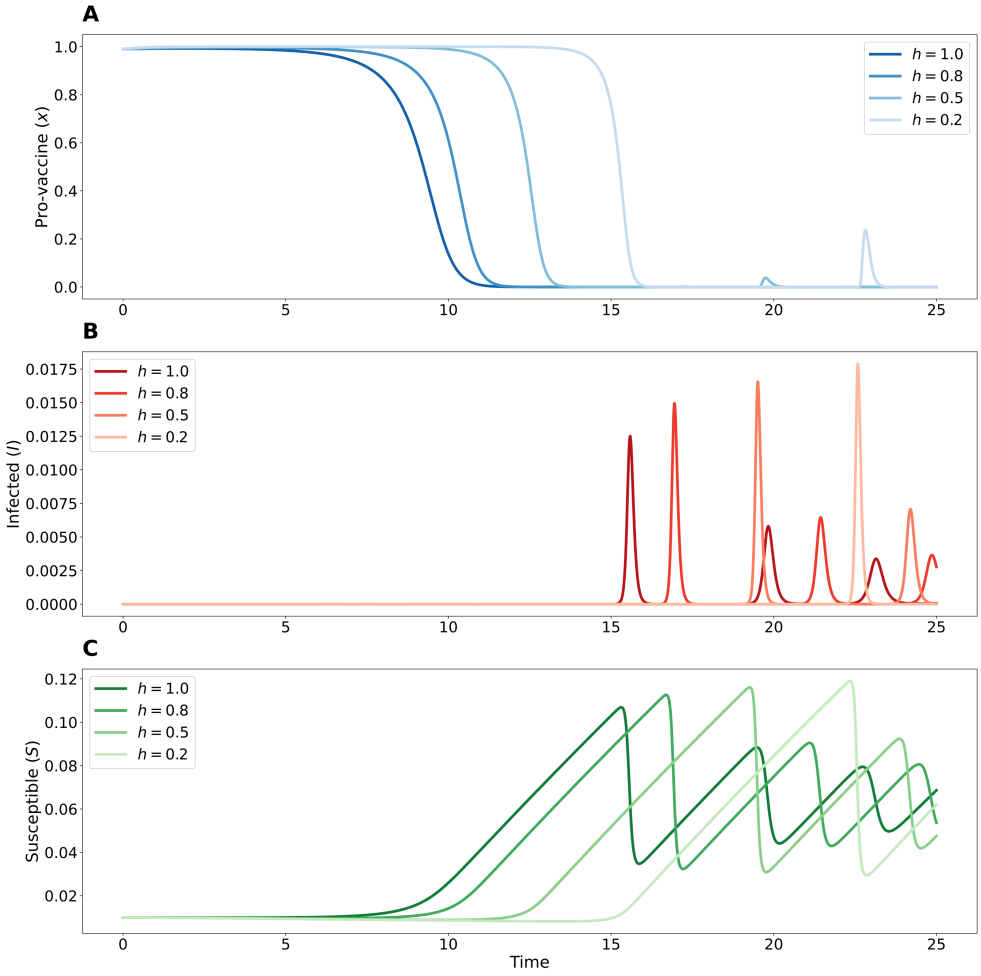


Figure 2.3: Homophily leads to earlier disease outbreaks. Simulations of (A) pro-vaccine sentiment percentage in the total population, (B) infection prevalence, and (C) susceptible percentage with different homophily levels: $h = 1.0/0.8/0.5/0.2$ (dark to light colors).

the critical transition.

We explored scenarios where the vaccine risk exhibited a linear increase. Notably, lower levels of homophily correspond to (i) delayed vaccine uptake collapse and subsequent disease outbreaks, as the pro-vaccine strategies enjoy wider dissemination (Figure 2.3A and B); (ii) a larger outbreak size due to the accumulation of susceptible individuals (Figure 2.3B and C). Moreover, lower homophily levels are associated with (iii) accelerated vaccine uptake collapse, evidenced by a steeper drop in response to the widespread anti-

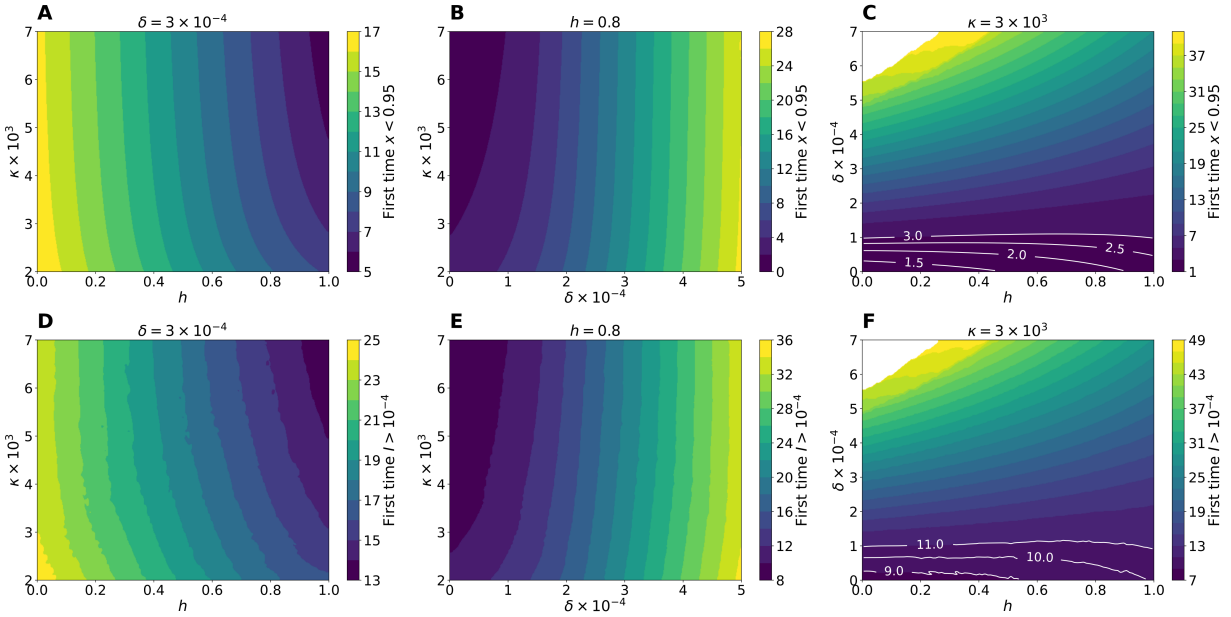


Figure 2.4: The first time when x drops below 95% in (A) h–κ (B) δ–κ, and (C) h–δ parameter planes. The first time when I increases above 10^{-4} in (D) h–κ (E) δ–κ, and (F) h–δ parameter planes. Model converges to the disease-free unanimously pro-vaccine equilibrium in the white regions in (C, F).

vaccine strategies with increasing vaccine risk (Figure 2.3A).

We also varied values of h (homophily level), δ (social norm strength), and κ (sampling rate), and examined their effects on two crucial time points: the first time when the population's pro-vaccine percentage x drops below 95%, the approximate herd immunity threshold for measles [67] (Figure 2.4A–C), and the time of the first epidemic when the infectious percentage I increases above 10^{-4} (Figure 2.4D–F). Keeping h and δ fixed, an increased κ precipitates an earlier collapse of vaccine uptake and disease outbreak (Figure 2.4A and D); maintaining h and κ constant, a higher δ leads to a delayed decline in vaccine uptake and subsequent disease outbreak (Figure 2.4B and E). Notably, our simulations reveal that with a small δ, higher homophily levels result in a later collapse in vaccine uptake as well as the disease outbreak (Figure 2.4C and F). This phenomenon arises because, in scenarios of low δ, the pro-vaccine social norm lacks the strength to promote such a strategy, thereby leading vaccinators to imitate non-vaccinators. However, homophily disrupts this trend, highlighting the advantageous role of reduced social interactions in such contexts.

Insights into sensitivity are also revealed from the parameter plane plots in Figure 2.4

(for the elasticity coefficient plot, please see Appendix A.3). By generating colormap indices based on discrete intervals, we observe that the regions in Figure 2.4A, B, D, and E are divided nearly horizontally, suggesting that variations in κ have minimal impact on the model outcomes. Additionally, we find that the model is converging towards the disease-free and unanimously pro-vaccine equilibrium, even in scenarios with linearly increasing vaccine risk. This trend only occurs when the homophily level is low and the social norm strength is high (Figure 2.4C and F).

2.3.4 Online events could significantly impact disease dynamics

Throughout the study, we assume that the population is originally dominated by pro-vaccine sentiments. We also believe the sequence of events shown in Figure 2.1 could possibly occur:

(i) A decline in the number of individuals holding pro-vaccine sentiments would lead to a corresponding decrease in the vaccination rate shortly. Meanwhile, the drop in pro-vaccine proportion (when the pro-vaccine faction is the majority) would intensify discussions about vaccines, as the proportions of individuals with opposite sentiments approach parity.

(ii) When the vaccination rate falls below a certain threshold, along with specific trigger events, such as the introduction of imported cases, the potential for a disease outbreak becomes a significant concern.

We posite that a prospective disease outbreak and the observed intense online discussions should not be regarded as unrelated events. In essence, when vaccination becomes a trending topic on social media, it serves as an indicator that the pro-vaccine percentage is on the decline, thereby signaling a potential decrease in the vaccination rate, a critical factor contributing to the likelihood of a future outbreak. It is also worth noting that the relationship between online discussions and sentiment change is a two-way street; not only do online discussions reflect the change of sentiments, but they also influence people's sentiments at the same time. Posts by official government and media accounts are often perceived as authoritative and more influential; accounts of celebrities tend to amass a larger number of followers, so any expressions of opinion from such figures can potentially trigger a substantial volume of re-posts. Despite their low probability of occurrence, such extreme events typically polarize public opinions, leading to a rapid surge in the percentage of one sentiment (and a drop in the other). Hence, in a simulation, the pro-vaccine

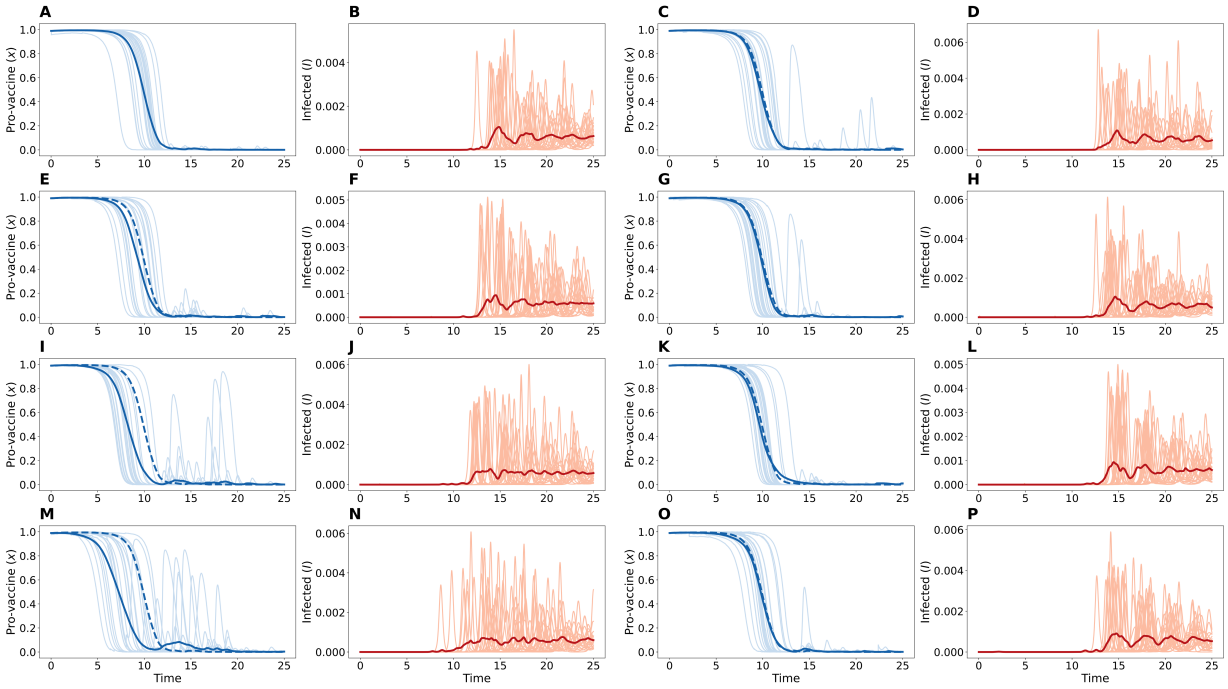


Figure 2.5: Possible realizations of the stochastic model. (A, E, I, M) Pro-vaccine percentage with increasing magnitudes of noise σ_{x_1} (with $N_1/N = 5\%$); (B, F, J, N) corresponding infection percentage. (C, G, K, O) Pro-vaccine percentage with different proportions of social media users N_1/N: (C)10%, (G) 20%, (K) 40% and (O) 60%; (D, H, L, P) corresponding infection percentage. Dark blue/red lines are the average of 100 simulations, 20 of which are shown using light lines. Dashed blue lines in (C, E, G, I, K, M, O) are the solid blue line in (A), the standard case.

percentage would undergo a steep increase (or decrease) when such events occurred in the absence of an outbreak.

Several realizations of the stochastic model are presented in Figure 2.5, in which we postulate an increasing vaccine risk, consequently expecting a gradual decline in the pro-vaccine percentage. We investigated the impact of social media activities on disease dynamics, where we assumed that social media behaviors can vary across different populations in two fundamental ways. Firstly, how people engage in online social activities can vary. To explore those variations, we used different noise magnitudes of σ_{x_1}. Basically, a larger σ_{x_1} signifies higher social engagement, resulting in accelerated opinion propagation and a higher possibility that extreme events would occur. Secondly, the proportion of social

media users can be different, which can be modeled by changing the ratio of social media users and non-users in a population while keeping the total population constant.

In Figure 2.5, we present the average curves of the overall pro-vaccine percentage derived from 100 realizations with various parameter sets. Notably, we see a significant shift in the critical transition point when we amplify the noise magnitudes σ_{x_1} (Figure 2.5A, E, I, and M). This shift suggests that heightened social engagement and an increased likelihood of extreme events could potentially accelerate the onset of vaccine uptake collapse. We also observe that disease outbreaks are more likely to occur earlier in scenarios with higher social engagement (Figure 2.5B, F, J, and N). Conversely, the proportion of social media users in a population does not greatly influence disease dynamics. While a larger proportion of social media users does make it easier for the overall pro-vaccine percentage to be affected by online events, we find that the timing of the critical transition displays no obvious difference (Figure 2.5C, D, G, H, K, L, O, and P).

2.3.5 Model explains the "early discussions" observed in social media data

A recent study has unveiled compelling evidence of CSD from Twitter data related to measles and the MMR vaccine. Researchers observed significantly increasing variance, lag-1 autocorrelation, and coefficient of variation in the geocoded time-series data. Remarkably, these patterns emerged years prior to the Disneyland measles outbreak in California in 2014 [175]. Our investigation extends beyond the California outbreak, encompassing major measles outbreaks in the United States, the United Kingdom, and Canada spanning the years 2012 to 2019. Similar patterns are found in the preceding Twitter data of each outbreak. In some instances, these patterns are less conspicuous due to the limited number of Twitter posts available for analysis.

We present the results focusing on three measles outbreaks in the selected countries for which we have successfully collected decent quantities of geocoded Twitter posts (Figure 2.6E–G, for details on processing the data, please see Section 2.5.2). We also plot the variance, lag-1 autocorrelation, and coefficient of variation of the selected time-series Twitter data (Figure 2.6I–K, M–O, and Q–S). In the California dataset (Figure 2.6A, E, I, M, and Q), a significant pattern emerges, showcasing a continuous increase in variance starting from mid-2013, approximately 1.5 years before the outbreak. Meanwhile, lag-1 autocorrelation and the coefficient of variation display an overall upward trajectory with some oscillations. An abnormal surge in posts results in an instant boost in all three indicators in the New York dataset in May 2017, followed by a rapid decline in December 2017

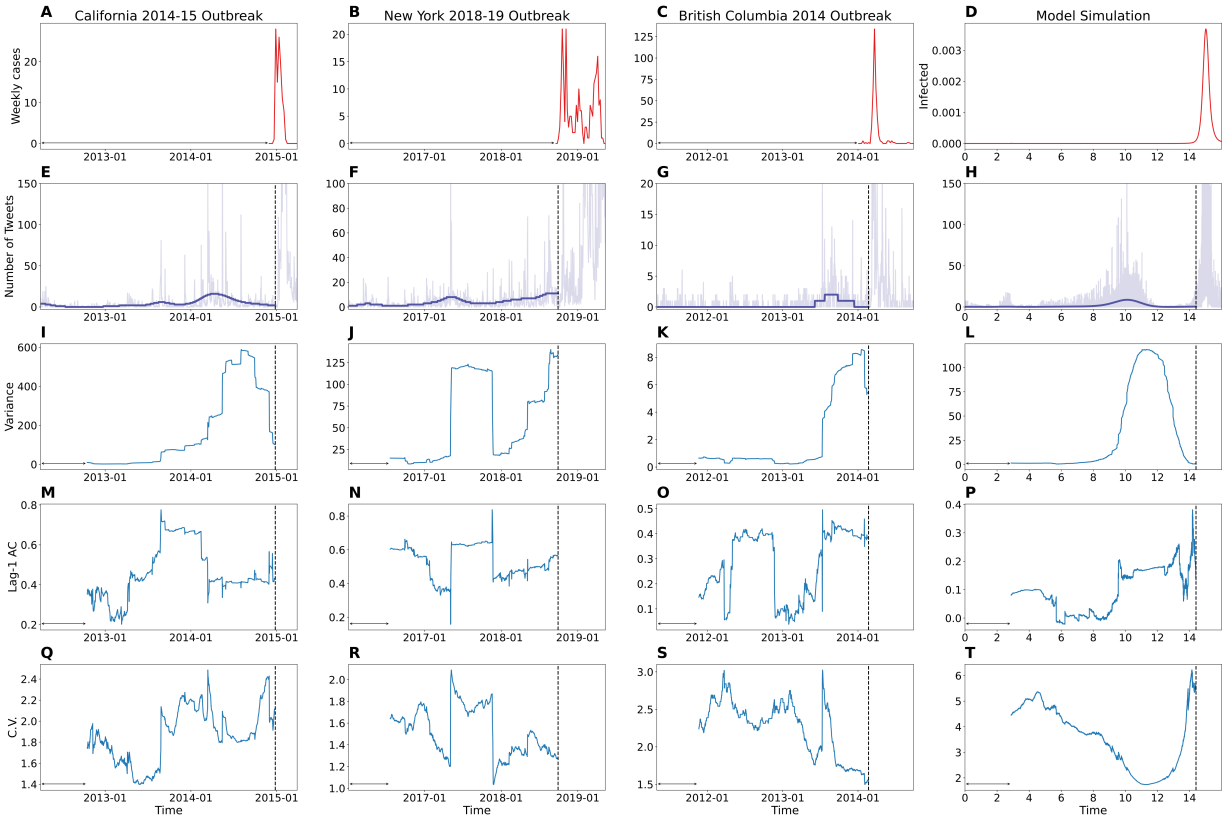


Figure 2.6: Measles cases/infection percentage, number of Twitter posts, and selected EWS. (A, E, I, M, Q) 2014–15 outbreak in California, (B, F, J, N, R) 2018–19 outbreak in New York State, (C, G, K, O, S) 2014 outbreak in British Columbia, Canada, and (D, H, L, P, T) one realization of the stochastic model. Selected EWS include (I–L) variance, (M–P) lag-1 autocorrelation, and (Q–T) coefficient of variation. The dark solid lines in (E–H) are smoothed curves using a Gaussian kernel. Vertical dashed lines in (E–T) are the starting times of the outbreaks.

(Figure 2.6B, F, J, N, and R). This shift is attributed to a substantial volume of retweets generated by a news report titled "Anti-vaccine activists spark a state's worst measles outbreak in decades" in *The Washington Post* [214], related to an outbreak in Minnesota in early 2017. Despite this transient disturbance, both variance and lag-1 autocorrelation continue to exhibit an overall upward trend, while the coefficient of variation displays a declining trend. In the case of British Columbia, Canada, the quantity of collected Twitter posts is relatively lower, resulting in less informative results (Figure 2.6C, G, K, O, and

S). Nevertheless, consistent patterns of variance and lag-1 autocorrelation can be observed, both of which increase years prior to the actual outbreak. Interestingly, the coefficient of variation displays an opposing behavior in its trajectory.

According to Figure 2.1, we consider a drop in vaccine coverage as a necessary condition for potential disease outbreaks, which is caused by a drop in the overall pro-vaccine percentage. The online pro-vaccine percentage, x_1, influences the overall percentage via Eq. (2.11). Due to the fact that the function $f(x) = x(1-x)$ is decreasing when $x > 0.5$, when x_1 falls from somewhere near 1, the portion $cx_1(1-x_1)$ in Eq. (2.12) increases. With the absence of disease incidence so that the term τI remains small, the number of social media posts is expected to increase.

In Figure 2.6H, we present the estimated numbers of social media posts prior to the first major outbreak in one possible realization of our stochastic model (Figure 2.5A). Initially, the number of social media posts remains relatively low. As the pro-vaccine percentage begins to decline, the number of posts exhibits increased variability, with more days witnessing a large number of posts. The smoothed curve also shows an increasing trend. Attributed to the linearly increasing vaccine risk parameter ω, the vaccine uptake collapses after passing the tipping point, subsequently triggering an epidemic. The number of posts decreases and becomes less variable since anti-vaccine sentiments gradually take the dominant position.

Indeed, similar patterns are observed in real-world data, especially in the California and British Columbia datasets shown in Figure 2.6E and G. However, it is worth noting that anti-vaccine opinions rarely dominate vaccine-related discussions. We would contend that these observed patterns are a reflection of people's reactions to the decline in the pro-vaccine percentage. Rather than x_1 diminishing from 1 to 0 and $1-x_1$ increasing from 0 to 1, x_1 exhibits a resurgence back to 1 following its initial decline from 1, while $1-x_1$ drops back to 0 after an initial increase from 0.

Compared with the EWS observed from the empirical data, the realization of our model obtains high resemblance on variance, which continuously increases years before the outbreak, followed by a subsequent decline (Figure 2.6L). The lag-1 autocorrelation of the model's realizations displays slight fluctuations but maintains an overall increasing trend, much like the trend observed in the real-world data (Figure 2.6P). In spite of the similar increasing trend, specific extreme events (Figure 2.6N) and large-scale discussions within a short period of time (Figure 2.6O) cause large jumps on lag-1 autocorrelation in the real-world data.

A key distinction arises in the coefficient of variation. The simulated data showcase a consistent pattern where the coefficient of variation decreases as the system approaches the

tipping point and subsequently increases as it transitions to a new regime (Figure 2.6T). This pattern is completely opposite to the behavior of variance, which is absent in the real-world data (Figure 2.6Q–S).

2.4 Discussion

In this article, we present a coupled behavior-disease model incorporating imitation processes between two distinct social groups (social media users and non-users) within an SIR framework. This innovative model enhances our understanding of vaccination dynamics concerning childhood diseases such as measles. Our findings indicate that homophily significantly accelerates the collapse of vaccine uptake, leading to an earlier potential outbreak. The stochastic model, driven by an α-stable Lévy process, effectively captures extreme events on social media, emphasizing the pivotal role of randomness in social media in influencing disease dynamics. Additionally, our model offers insight into the temporal dynamics of heated social media discussions preceding disease outbreaks, aligning with the EWS observed in a recent study [175].

Homophily disrupts intergroup learning [159], irrespective of the nature of the strategies involved. This implies that, regarding vaccination, the spread of both pro-vaccine and anti-vaccine strategies is prevented due to homophily. Essentially, in an open environment where pro-vaccine strategies gain widespread adoption, anti-vaccine strategies are easily spread at the same time. Given our focus on achieving a unanimously pro-vaccine status, the social norm favoring the adoption of the pro-vaccine strategies mitigates the influence of the anti-vaccine strategies. However, in scenarios where the vaccine risk is exceptionally high, leading to the prevalence of anti-vaccine norms, a lower level of homophily exacerbates the collapse of vaccine uptake. Our model also reveals that when the homophily level is very low and the social norm strength is very high, the system can still converge to the disease-free and unanimously pro-vaccine equilibrium even if the vaccine risk is continuously increasing.

Online discussions not only serve as reflections of individuals' attitudes toward vaccination but also have a substantial influence on their choices regarding vaccination strategies. Our stochastic model highlights the role of stochasticity as a core determinant impacting disease dynamics. In a highly active online environment where the occurrence of extreme events is more probable, strategy shifts among individuals are more likely to happen. Consequently, the overall pro-vaccine percentage easily drops, forcing the system to approach the tipping point. Our results indicate that the proportion of social media users is not a predominant factor. Its influence is integrated into the system through the computation of the overall pro-vaccine percentage, which, in turn, determines the proportion of susceptible

and recovered newborns. As the engagement in online discussions increases, the overall pro-vaccine percentage converges towards the online pro-vaccine percentage. While it is true that the online pro-vaccine percentage experiences an earlier drop compared to the offline percentage, the magnitude of this difference is relatively small.

We expand upon existing research focused on observing EWS in social media data [175], uncovering evidence across diverse datasets spanning the United States, the United Kingdom, and Canada. Despite these valuable insights, the evidence remains limited for several reasons. Firstly, measles is not a widespread infectious disease thanks to the success of vaccination programs, resulting in infrequent and relatively small-scale outbreaks. Secondly, the act of posting Tweets to express sentiments about vaccination may not be a common practice for the majority of social media users. In regions with smaller populations, data collected from social media may be highly biased. Thirdly, biases are introduced during data manipulation processes, involving steps such as text cleaning, geocoding, and sentiment analysis, which leads to the exclusion of a substantial portion of the raw dataset. Lastly, the challenge of identifying and isolating overlapping events on social media is almost impossible. Events occurring elsewhere can impact online activities in the target area, contributing to potential confounding factors in the interpretation of results.

We attempt to establish a connection between infectious disease dynamics and online social media activities through mathematical models. Given the practical difficulty of obtaining a substantial volume of real-world social media data reflecting sentiments about vaccines, generating artificial data using mathematical models becomes essential for potential future research work. Recent advancements show that deep learning models are capable of detecting EWS from time-series data [38]. This capability opens the possibility of employing such models for the direct detection of EWS from social media data. However, generating realistic social media data remains a challenge due to the complex nature of human behaviors in drafting social media posts. For instance, the frequency of days with no posts collected through social media far exceeds the days with zero posts obtained from the model. Moreover, the realistic validity of our model's realizations remains unverifiable.

Our model assumes a linearly increasing vaccine risk when generating theoretical data. While this assumption provides a simplified model, it may not accurately reflect real-world dynamics, where society typically reacts to declines in vaccination rates. As a result, the pro-vaccine proportion is unlikely to drop to zero as observed in our stochastic simulation. A more realistic representation of vaccine risk could involve its correlation with the prevailing pro-vaccine percentage or the introduction of a piecewise function that declines after an increase [21]. However, such functions often introduce complexity, expanding the parameter space to higher dimensions, which may not be ideal in the context of model selection. It is crucial to acknowledge that our model's accuracy diminishes over the entire

time series. Yet, it remains reasonably precise, particularly in the segment preceding the tipping point, as suggested by some key indicators.

We categorize the population into two distinct social groups, namely social media users and non-users, and show the imitation dynamics between them. Nevertheless, the method of population division is not unique. The variables x_1 and x_2 can denote the pro-vaccine percentage of any two separate and non-overlapping groups. The sampling rate (κ) and social norm strength (δ) could generally differ for each group. However, we choose to use the same parameters for simplicity, as our focus is on the effect of homophily. In the stochastic model, such differences are represented by different scaling factors of the additive noises. Furthermore, the model can be extended to include more than two disjoint groups. For instance, the variables x_i for $i \in \{1, 2, \ldots, n\}$ could represent individuals in various age groups, such as children, youth, adults, and seniors.

Throughout the article, we assume that homophily only affects opinion exchange, and thus analyze the disease dynamics over a homogeneous population. Its impact on direct social interactions should also be considered, necessitating the use of a patch model [102] or a multi-layer social network [104]. Other factors such as vaccine effectiveness [50] and individual decision-making processes (e.g. commitment to vaccinate) could be discussed for a more comprehensive understanding of the factors influencing vaccination behaviors.

Our investigations unveil the influence of homophily and online randomness on disease transmission. Moreover, our model effectively explains early discussions preceding various measles outbreaks. Its accuracy near tipping points could be beneficial for future research grappling with insufficient social media datasets, offering a valuable tool for nuanced exploration.

2.5 Methods

2.5.1 Measles outbreaks and cases

An outbreak of a disease is characterized by a sudden surge in cases surpassing the normal expectancy in a limited geographic area [69]. As an extremely contagious disease, a measles outbreak is declared by the World Health Organization (WHO) when at least two laboratory-confirmed cases are temporally related [238]. Consequently, many reported measles outbreaks, often caused by imported cases in measles-endemic regions, exhibit relatively modest sizes. These short-lived outbreaks usually do not attract the widespread

Location	Starting time	Length	Total cases
British Columbia, Canada	January 2014	36 weeks	331
Ohio, United States	March 2014	121 days	383
California, United States	December 2014	42 days	125
Quebec, Canada	January 2015	72 days	159
London, United Kingdom	January 2016	12 months	283
Minnesota, United States	March 2017	58 days	65
London, United Kingdom	January 2018	9 months	211
New York, United States	October 2018	211 days	242
Washington, United States	December 2018	73 days	71

Table 2.2: Major measles outbreaks in the US, Canada and England, UK from 2014 to 2019

attention of the public. Table 2.2 highlights selected outbreaks that endured for at least six weeks, during which a minimum of 60 confirmed cases were reported.

Daily cases for four major measles outbreaks in the US (CA 2014, MN 2017, NY 2018, WA 2018) were collected from CDC *Morbidity and Mortality Weekly Report (MMWR)* [246, 107, 41, 156]. Daily cases during the 2014 Ohio outbreak were collected from an *MMWR* report [96] and an article in *The New England Journal of Medicine* [95]. Weekly cases for two outbreaks in Canada (BC 2014, QC 2015) were obtained from *Canada Communicable Disease Reports (CCDR)* [203, 202]. Before WHO declared measles elimination in the UK in 2017, regular outbreaks occurred. Measles transmission resurged in 2018. For England, two major outbreaks in London were selected to represent periods pre- and post-2017, as per *Health Protection Report* by Public Health England [186, 187].

2.5.2 Twitter data

For the social media dataset, we collected English Tweets containing the keywords "measles OR MMR" from January 1, 2011, to December 31, 2019, using the Twitter Application Programming Interface (API). The collected Tweets underwent text cleaning, which involved removing usernames, URL links, non-English characters, and extra whitespaces. In hashtags, the hash symbol # was removed, retaining the text for potential sentiment analysis. Emojis (pictograms embedded in text) were transferred to pure texts (demojized) using the Python package *emoji*. After filtering out Tweets lacking the specified keywords, the dataset comprised 4,545,965 valid Tweets.

Subsequently, we employed geocoding techniques using the Python package *geopy* (for

Tweets with coordinates) and a modified algorithm from the Python library *geodict* [231] (for location information present in text). For Tweets with coordinates, we extracted city, region (state for the US, province for Canada, county for England), and country information. Tweets with location information in text were processed by identifying valid location name templates like "city, region", "region, country", and "city, country". In cases where two cities with the same name existed in different regions, the one with a larger population was selected. Following geocoding, approximately 40% of the Tweets remained, yielding 1,214,641 Tweets in the US, 147,198 Tweets in Canada, and 137,119 Tweets in England.

We calculated the daily number of pro-vaccine Tweets and collected the most recent 1000 data points in California/New York/British Columbia before the start of the selected outbreaks (CA 2014, NY 2018, BC 2014). We then smoothed the time-series data using the Python package *ewstools* [36] with a Gaussian kernel at a bandwidth of 10%. The variance and lag-1 autocorrelation (correlation between values that are one time period apart, i.e. $Corr(\text{data}_t, \text{data}_{t-1})$) of the residual time series, and the coefficient of variation of the original time series were calculated using a rolling window with a width of 20% of the data.

2.5.3 Sentiment analysis

Each cleaned and geocoded Tweet was classified into one of three categories: pro-vaccine, anti-vaccine, or neutral, based on its sentiment towards vaccination. Pro-vaccine Tweets exhibited a positive stance, including supportive actions or sentiments. Conversely, anti-vaccine Tweets conveyed a negative attitude. All other Tweets, including those with ambiguous attitudes or objective content like news reports, were classified as neutral.

We employed COVID-Twitter-BERT [168], a transformer-based pre-trained natural language processing (NLP) model, for Tweet classification. While optimized for COVID-19-related Tweets, the model demonstrates proficiency in classifying other disease-related content after fine-tuning with a small relative dataset. Experimenting with different ratios of pro-vaccine, anti-vaccine, and neutral Tweets for fine-tuning, the optimal performance with the best mean precision and recall scores occurred with a ratio of approximately $2:1:3$. We randomly selected 1,200 pro-vaccine, 600 anti-vaccine, and 1,800 neutral Tweets from the Vaccine Sentiment Dataset [167] for fine-tuning, using a learning rate of 2×10^{-5} across 5 epochs. The model achieved precision scores of [71.7%, 89.7%, 86.7%] and recall scores of [94.3%, 74.3%, 74.3%] for pro-vaccine, anti-vaccine, and neutral Tweets, respectively, over an equal number of instances for each sentiment category. Applying

the fine-tuned model on all geocoded Tweets, we obtained 608,983 pro-vaccine and 54,425 anti-vaccine Tweets.

Chapter 3

Forecasting infectious disease outbreak risks from vaccine sentiments on social media: A data-driven dynamical systems approach

This chapter is based on the paper: Z. He and C. T. Bauch. Forecasting infectious disease outbreak risks from vaccine sentiments on social media: A data-driven dynamical systems approach. *Mathematical Biosciences and Engineering*, 2025, 22(10):2761–2779.

3.1 Abstract

Early warning signals are vital in predicting critical transitions in complex dynamical systems. For behavioral epidemiology systems in particular, this includes shifts in vaccine sentiments that may precede disease outbreaks. Conventional statistical indicators, such as variance and lag-1 autocorrelation, often struggle in noisy environments and may fail in real-world scenarios. In this study, we leveraged universal signals of critical slowing down to train deep learning classifiers, specifically using long short-term memory (LSTM) and residual neural network (ResNet) architectures, for detecting early warning signals in disease-related social media time series. These classifiers were trained on simulated data from a stochastic coupled behavior-disease model with additive Lévy noise, a non-Gaussian noise that better reflects the heavy-tailed nature of real-world fluctuations. Our results show that these classifiers consistently outperform conventional indicators in both sensitivity and specificity on theoretical data while delivering quantitatively clear results that are easier to interpret on empirical data. Integrating deep learning with real-time social media monitoring offers a powerful tool for preventing disease outbreaks through proactive public health interventions.

3.2 Introduction

Before the development of vaccines, infectious diseases like measles were widespread and often resulted in significant morbidity and mortality, particularly among children [43]. The introduction of vaccines drastically reduced the incidence of these diseases, saving countless lives and alleviating the burden on healthcare systems [44]. Vaccine coverage is crucial not only for individual protection but also for establishing herd immunity, which prevents the spread of diseases within communities, including to those who cannot be vaccinated due to medical reasons or choose not to get vaccinated because of personal interests [7, 67]. When the vaccination rate drops below the herd immunity threshold, the risk of disease transmission increases, and any trigger event, such as an imported case, can potentially lead to a large outbreak. Between 2000 and 2013, California saw a concerning decline in the rate of measles, mumps, and rubella (MMR) vaccination [40, 134]. This drop in vaccination coverage resulted in the loss of herd immunity, culminating in a significant measles outbreak at Disneyland, in California, in late 2014 [246, 236, 71]. This event underscored the critical need for maintaining high vaccination rates to prevent the resurgence of preventable diseases.

Monitoring vaccination rates is crucial. In the United States, the Centers for Disease

Control and Prevention (CDC) oversees vaccination data collection through the National Immunization Survey (NIS) [42]. This annual survey gathers information via telephone from households with children and adolescents, cross-referencing data with healthcare providers' records. In Canada, the Public Health Agency of Canada (PHAC) collaborates with Statistics Canada to collect national vaccination data, conducting the Childhood National Immunization Coverage Survey (cNICS) every two years since 1994 [171]. In England, Public Health England (PHE) manages childhood vaccination data collection through the Cover of Vaccination Evaluated Rapidly (COVER) program, which operates on an annual basis [219]. However, annual or biennial vaccination data collection has notable limitations concerning bias and timeliness [60]. Response bias can significantly affect the reliability of such data, as those who choose to participate in surveys may not accurately represent the broader population, potentially skewing the results. Timeliness is another critical issue. While these surveys provide periodic updates, they may not capture rapid changes or emerging trends in vaccination rates between survey periods. This gap can delay the identification of vaccination coverage issues and the implementation of timely public health interventions. Consequently, the data might not reflect the current state of vaccination coverage, hindering efforts to address potential outbreaks promptly.

Social media data offers a high degree of timeliness, as it is generated and updated instantaneously, providing immediate insights into public behaviors, sentiments, and discussions. This real-time nature makes social media an attractive tool for infectious disease surveillance [2, 234, 184]. Public health authorities can quickly detect emerging health concerns [8] and monitor changes in public attitudes toward vaccinations [194, 3] by analyzing social media activities, such as posts, shares, and comments. Recent advancements have also shown significant promise in incorporating social media into infectious disease modeling [209]. Social media plays a crucial role in altering people's behaviors by raising awareness of the disease, which generally helps to slow its spread. For example, when the number of infections increases and is reported on social media, susceptible individuals become more cautious and take protective measures, leading to a decrease in the transmission rate between susceptible and infected individuals, which is often modeled by using an exponentially decreasing function [59, 242, 162]. Some models also introduce separate compartments for aware and unaware individuals and model the transitions among them [1].

As researchers explore ways to incorporate social media into mathematical models, the potential for enhancing early warning systems becomes increasingly promising [210]. Critical transitions and early warning signals (EWS) are fundamental concepts in understanding and forecasting sudden changes in complex systems [197, 63]. Critical transitions occur at moments when a dynamical system reaches a bifurcation point, leading to sudden

and often irreversible shifts. EWS, on the other hand, typically precede these transitions and can provide notice of impending changes, potentially enabling the mitigation of adverse outcomes. Specific EWS, such as increased variance and lag-1 autocorrelation, have demonstrated reliability in indicating the emergence of diseases in epidemic models [32, 33].

Deep learning (DL) excels at time series classification due to its ability to learn complex patterns within sequential data [230, 123]. Research has demonstrated that deep learning models can predict different types of bifurcations in time series data for both continuous- and discrete-time models in physiology, economics, and ecology [38, 37, 76]. These deep learning models can recognize subtle changes in the time series data that precede critical transitions by learning generic features of bifurcations from large simulated training datasets. This capability is particularly useful in the context of disease outbreaks, where early detection of critical transitions can significantly improve response times and outcomes.

Certain EWS (e.g., variance, lag-1 autocorrelation) have been observed in vaccine-related Twitter data and Google trends prior to major measles outbreaks [175], indicating that changes in public vaccine sentiments can precede epidemiological events. However, traditional EWS methods typically assume white noise [146], whereas social media dynamics often follow non-Gaussian noise due to occasional high-impact events, such as viral posts from celebrities or influencers. These rare but large fluctuations create heavy-tailed distributions that can distort standard EWS metrics. To explain the phenomenon of "early discussion" and also better model the extreme events on social media, researchers have developed a stochastic coupled behavior-disease model with additive Lévy noise that links social media activity with disease dynamics [109]. This model suggests that heightened online discussion about vaccines and infectious diseases can serve as a signal to shifts in vaccine uptake and disease transmission. Building on these findings, there is promising potential for using deep learning models to detect EWS for vaccine uptake collapse in social media data. These models can learn to recognize patterns suggesting declining vaccination rates by analyzing large datasets, thereby providing timely alerts and enabling public health interventions to prevent outbreaks.

The growing field of behavioral epidemiology emphasizes the importance of incorporating human behavior into the modeling and control of infectious diseases [93, 152, 229, 223]. For example, vaccination as an effective measure of preventing infectious diseases has been extensively studied using evolutionary game theory to account for individual decision-making [20, 190, 94, 73, 204, 50, 48]. Over time, this line of research has considered factors such as vaccine scares [21, 205], social networks [89, 169], social norms [173], and homophily [109], leading to more realistic models of vaccination dynamics. These increasingly complex models not only capture key behavioral mechanisms but also offer an

opportunity: they can be used to generate synthetic data for training deep learning models. This idea, previously suggested by [14], is extended in this study to contribute new tools for the preparedness and mitigation of potential outbreaks.

In this article, we present a behavioral epidemiological method for using deep learning classifiers to provide EWS for potential vaccine uptake collapse. We train these classifiers using simulated data from a stochastic coupled behavior-disease model based on [109] and test their performance on real-world Twitter data collected before several major measles outbreaks in North America. We evaluate the classifiers' performance relative to traditional EWS indicators such as variance and lag-1 autocorrelation. In addition, we vary the noise amplitude to assess the robustness of these EWS methods.

3.3 Methods

3.3.1 Deep learning model framework

We used two deep learning algorithms. The first was a CNN-LSTM framework similar to [38, 37, 76], consisting of a 1D convolutional layer with dropout, a max pooling layer, and two LSTM layers with dropout, followed by a dense layer (Figure 3.1A).

The second algorithm applied ResNet [108], as inspired by [230] (Figure 3.1B). Each residual block consists of three convolutional layers with progressively smaller kernel sizes, enabling the model to capture different levels of abstraction. These layers are followed by batch normalization, ReLU activation, and dropout layers to prevent overfitting. The input to the block is added directly to the output, allowing the model to retain information from earlier layers. The overall architecture includes three such residual blocks. The model concludes with global average pooling and a dense layer, which produces the final output using a sigmoid activation.

Both models were compiled using the Adam optimizer with binary cross-entropy loss and trained for 300 epochs at a learning rate of 0.0005 and a batch size of 1024. All code was written in Python using TensorFlow 2.12.0 and executed in Jupyter Notebook. Notably, the ResNet model contains approximately 20 times more trainable parameters than the CNN-LSTM model and requires a much longer time to train. The CNN-LSTM model achieved precision scores of 86.3% and 99.3% and recall scores of 99.4% and 84.2% for positive and negative cases, respectively, using an equal number of instances for each category. In comparison, the ResNet model achieved precision scores of 96.2% and 80.3% and recall scores of 76.2% and 97.0% for positive and negative cases. The overall accuracy was 91.8% for the CNN-LSTM model and 86.6% for the ResNet model.

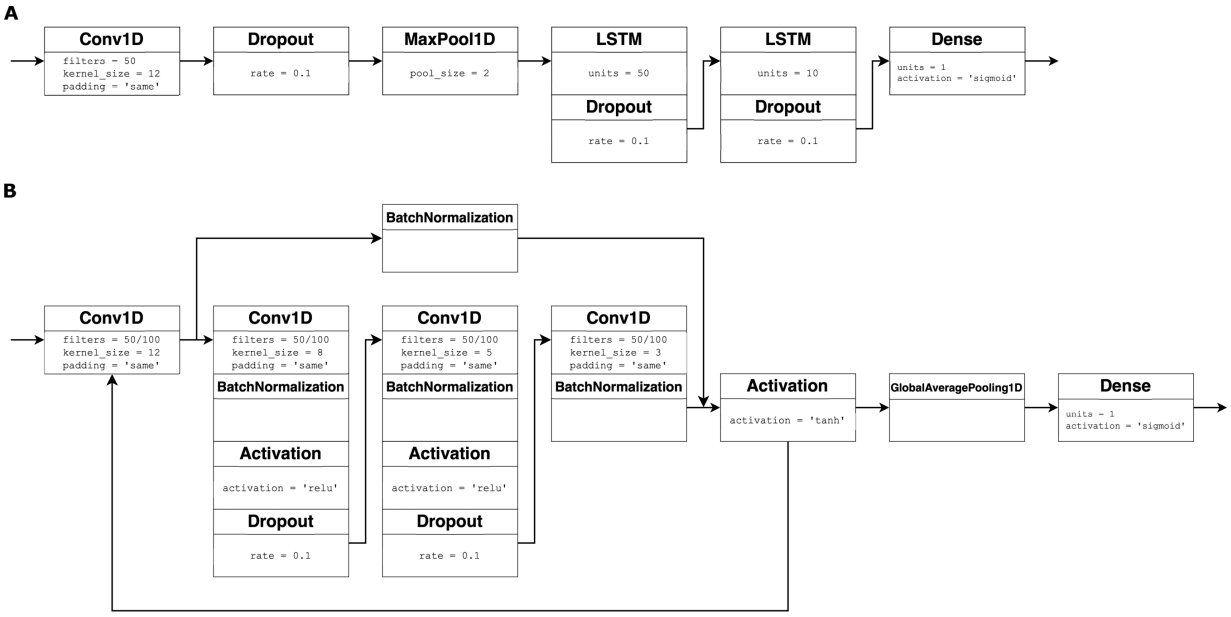


Figure 3.1: Two deep learning frameworks: (A) CNN-LSTM, (B) ResNet.

3.3.2 Theoretical data for training

We applied a coupled behavior-disease model that incorporates two different social groups (social media users and non-users) [109] with an exponentially decreasing media function [59, 242, 162]:

$$\frac{dS}{dt} = \mu(1-x) - \beta e^{-mI} SI - \mu S, \tag{3.1}$$

$$\frac{dI}{dt} = \beta e^{-mI} SI - \gamma I - \mu I, \tag{3.2}$$

$$\frac{dR}{dt} = \mu x + \gamma I - \mu R, \tag{3.3}$$

$$\begin{aligned}\frac{dx_1}{dt} = {} & \kappa x_1(1-x_1)[-\omega + I + \delta(2x_1 - 1 + (1-h)(2x_2-1))] \\ & + (1-h)\kappa\Big\{x_2(1-x_1)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\ & - (1-x_2)x_1 \max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\},\end{aligned} \tag{3.4}$$

$$\frac{dx_2}{dt} = \kappa x_2(1-x_2)[-\omega + I + \delta((1-h)(2x_1-1)+2x_2-1)]$$
$$+ (1-h)\kappa\Big\{x_1(1-x_2)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0]$$
$$- (1-x_1)x_2\max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\}, \tag{3.5}$$

where S, I, and R are the proportions of susceptible, infectious, and recovered individuals respectively with $S + I + R = 1$. x_1 (respectively, x_2) is the proportion of social media users (respectively, non-social media users) that are pro-vaccine. The overall pro-vaccine percentage in the population is the weighted average $x = (x_1N_1 + x_2N_2)/(N_1 + N_2)$. The total population size was set to $N = N_1 + N_2 = 200{,}000$, and the initial social media user ratio N_1/N was drawn from $\mathcal{U}[0.05, 0.4]$. We assumed that the number of social media users N_1 had an approximate annual growth rate of 5%. We also assumed that individuals sample others at a rate κ and adopt the sampled strategy with a probability proportional to the payoff difference, provided it is positive [20]. The basic reproduction number R_0 of measles was drawn from the uniform distribution $\mathcal{U}[12, 18]$ [7, 105], the recovery rate γ was drawn from $\mathcal{U}[365/22, 365/13]$ year^{-1} [173, 175], the natural birth/death rate μ was drawn from $\mathcal{U}[1/70, 1/50]$ year^{-1}, and the homophily level h was drawn from $\mathcal{U}[0.4, 0.8]$ [109]. For the remaining parameters, the transmission rate was calculated as $\beta = R_0(\gamma + \mu)$ [7], the sampling rate was set to $\kappa = 3000$ year^{-1} [175, 109], the social norm strength was set to $\delta = 3\times10^{-4}$ [173, 175, 109], and $m = 0.6$ quantified the impact of the media coverage, indicating a relatively weak behavioral response to increases in infection prevalence [59, 147]. This coupled model captures a population of social media users and non-users, where individuals imitate others' vaccination strategies both within and across social groups, but tend to interact more within their own group due to homophily [157].

The deterministic model was converted to a stochastic model driven by an α-stable Lévy noise [170],

$$d\mathbf{X} = \mathbf{f}(\mathbf{X}, t)dt + \boldsymbol{\sigma}dL, \tag{3.6}$$

where $\mathbf{X} = [S, I, x_1, x_2]$. $\mathbf{f}(\mathbf{X}, t)$ is a column vector of the right-hand sides of Eqs 3.1, 3.2, 3.4, and 3.5. $\boldsymbol{\sigma} = [\sigma_S, \sigma_I, \sigma_{x_1}, \sigma_{x_2}] = [10^{-3}, 2\times10^{-6}, 4\times10^{-3}, 10^{-3}]$ is a column vector of scaling factors controlling the magnitude of noise for S, I, x_1, and x_2. dL is the additive α-stable Lévy process. To preserve the total population in the stochastic model, we impose the constraint $\sigma_R = -\sigma_S - \sigma_I$, which ensures that the Lévy noise terms do not alter the overall population size. The number of pro-vaccine social media posts was estimated by

$$T_p = k[cx_1(1-x_1) + \tau I] + \xi. \tag{3.7}$$

Here, the scaling factor $\lambda = 0.2$, while c and τ were drawn from $\mathcal{U}[10, 200]$ and $\mathcal{U}[50{,}000, 100{,}000]$, respectively. ξ is an additive noise that follows a skewed fat-tailed stable distribution, with

an expected value of $(1-\lambda)[cx_1(1-x_1)+\tau I]$. The stochastic model was simulated using a self-developed Monte-Carlo algorithm. The additive Lévy noise term dL was generated using the *levy_stable* function from the Python module *scipy.stats* with parameters $\alpha = 1.5$ and $\beta = 0$; extremely large jumps with magnitude larger than 200 were disregarded. The stochastic term ξ was also generated using *levy_stable* with parameters dynamically adjusted based on the baseline level $cx_1(1-x_1)+\tau I$. To reflect the variability in user engagement, T_p was further perturbed using a probabilistic filter that randomly reduced its value based on posting likelihood. Finally, all solutions to Eq 3.6 were clipped to lie within the physical bounds $[0, 1]$ [109].

We modeled the bifurcation parameter, the relative vaccine risk ω, as a linearly increasing function $\omega = (2 + 0.5t) \times 10^{-4}$. A critical transition was defined heuristically when the deviation of x from the equilibrium was greater than 20 times the noise amplitude $\sigma_x = (\sigma_{x_1} N_1 + \sigma_{x_2} N_2)/(N_1 + N_2)$, indicating a likely escape from the basin of attraction [37]. Such a transition should occur before x drops below 95%, the approximate herd immunity threshold for measles [67]. Upon detecting a critical transition, the 500 data points of T_p immediately preceding the transition were classified as a "positive" series; if no critical transition was observed, the time series was disregarded. A "negative" (or "neutral") series was generated by setting ω just below the estimated bifurcation point $\delta(2-h)$ [109] and collecting the final 500 data points. Each series was then smoothed using the Python package *ewstools* [36] with a Gaussian kernel at a bandwidth of 100 data points. The residual and the smoothed time series together form a two-dimensional sample. A balanced set of 10,000 samples of each type was used for training.

3.3.3 Empirical data for testing

We collected English Tweets containing the keywords "measles" or "MMR" from January 1, 2011, to December 31, 2019, using the Twitter API. After cleaning the text (e.g., removing usernames, URLs, non-English characters, and converting emojis to text), we used the Python package *geopy* and a modified version of the *geodict* library [231] to geocode location information. Deep learning has been widely used for sentiment analysis of vaccine-related content on social media [12, 6]. Here, the cleaned and geocoded Tweets were then classified as pro-vaccine, anti-vaccine, or neutral using COVID-Twitter-BERT [168]. Our analysis focused on four major measles outbreaks in the United States and Canada (California 2014–15, New York 2018–19, Washington 2018–19, and British Columbia 2014) [246, 156, 41, 203] as well as four countries with mandatory measles vaccination policies and low incidence rates between 2014 and 2018 (Singapore, Mexico, Argentina, and Brazil) [199, 154, 241, 240]. Pro-vaccine Tweets from the selected regions and time frames were smoothed using

ewstools with a Gaussian kernel at a bandwidth of 100 data points. The variance and lag-1 autocorrelation of the residual time series were then computed using a rolling window of 200 data points.

3.4 Results

3.4.1 Performance of EWS and DL on theoretical data

We simulate sample positive and negative (or neutral) time series with different noise amplitudes. Time series with small amplitudes exhibit minor oscillations in the pro-vaccine percentage (Figures 3.2A and C), while those with large amplitudes display more pronounced oscillations (Figures 3.2B and D). We monitor two conventional EWS, variance and lag-1 autocorrelation, of the time series data on pro-vaccine percentage (Figures 3.2E–L) and the number of pro-vaccine social media posts (Figures 3.3E–L). EWS indicators are expected to show increasing trends as systems approach critical transitions. Consistent with this, we observe strong increasing trends in variance for both the pro-vaccine percentage (Figures 3.2E and F) and the number of pro-vaccine social media posts in positive samples (Figures 3.3E and F). However, lag-1 autocorrelation does not exhibit a similar upward trend (Figures 3.2I and J, Figures 3.3I and J).

The outputs of two deep learning classifiers applied to the number of pro-vaccine social media posts are also shown (Figures 3.3M–P). The probabilities assigned by the classifiers are smoothed using a Gaussian kernel to reveal an overall trend. For positive cases, both frameworks predict probabilities near zero for a long time before sharply increasing, effectively signaling the onset of a transition. In particular, LSTM transitions earlier than ResNet, indicating greater sensitivity to early signals, and a faster response. For negative cases, ResNet consistently predicts probabilities close to zero throughout the time, whereas LSTM produces several spikes and oscillations following sudden increases in the number of pro-vaccine posts.

The performance of these indicators is assessed using receiver operating characteristic (ROC) curves (Figure 3.4). The upward trends in variance and lag-1 autocorrelation are quantified by calculating the Kendall tau coefficient, while the final data point in the smoothed probabilities assigned by the deep learning classifiers represents their prediction outcome. We simulate a balanced set of 500 positive samples and 500 neutral samples across five different noise amplitudes. The statistics are generated over the entire pre-transition period, as well as on partial pre-transition periods up to 100, 200, and 400 time points

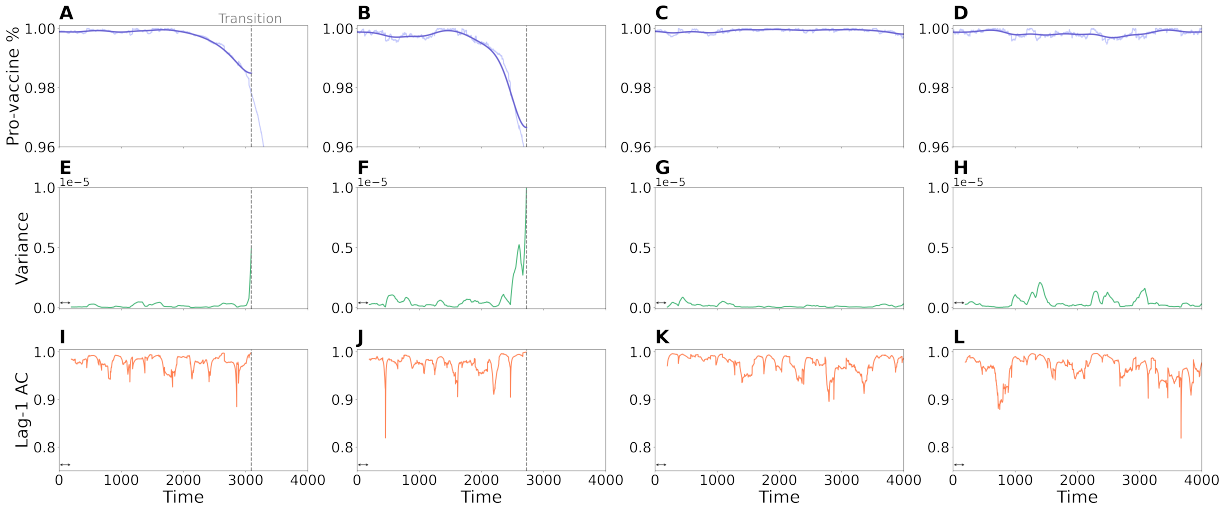


Figure 3.2: Trends in two conventional EWS (variance and lag-1 autocorrelation) applied to simulated time-series data. (A–D) Pro-vaccine percentage x in the entire population (light purple) and its smoothed trend (dark purple) for different cases: (A) a positive sample with small noise amplitudes, (B) a positive sample with large noise amplitudes, (C) a neutral sample with small noise amplitudes, and (D) a neutral sample with large noise amplitudes. (E–H) Variance of x. (I–L) Lag-1 autocorrelation of x. In (A, B), grey dashed lines mark the transition time.

before the transition. For neutral samples, the "pre-transition" period is defined as the entire time series, as no transition will occur. Using the area under the ROC curve (AUC) score, we find that variance on T_p is the best-performing conventional EWS, while lag-1 autocorrelation on T_p performs the worst (Figures 3.4E–H). Both indicators on x exhibit weak predictive ability (Figures 3.4A–D). Deep learning classifiers consistently outperform both variance and lag-1 autocorrelation on x and T_p at all cut-off times in the pre-transition period (Figures 3.4I–L). When evaluated over the entire pre-transition period, both classifiers achieve near-perfect performance, with ROC curves approaching the top-left corner and AUC scores close to 1, indicating high sensitivity and specificity. However, their performance gradually declines for both classifiers as the cut-off time moves further from the transition, with LSTM generally outperforming ResNet. This trend reverses at the 400-time-point cut-off, where LSTM's performance drops significantly, particularly for negative samples.

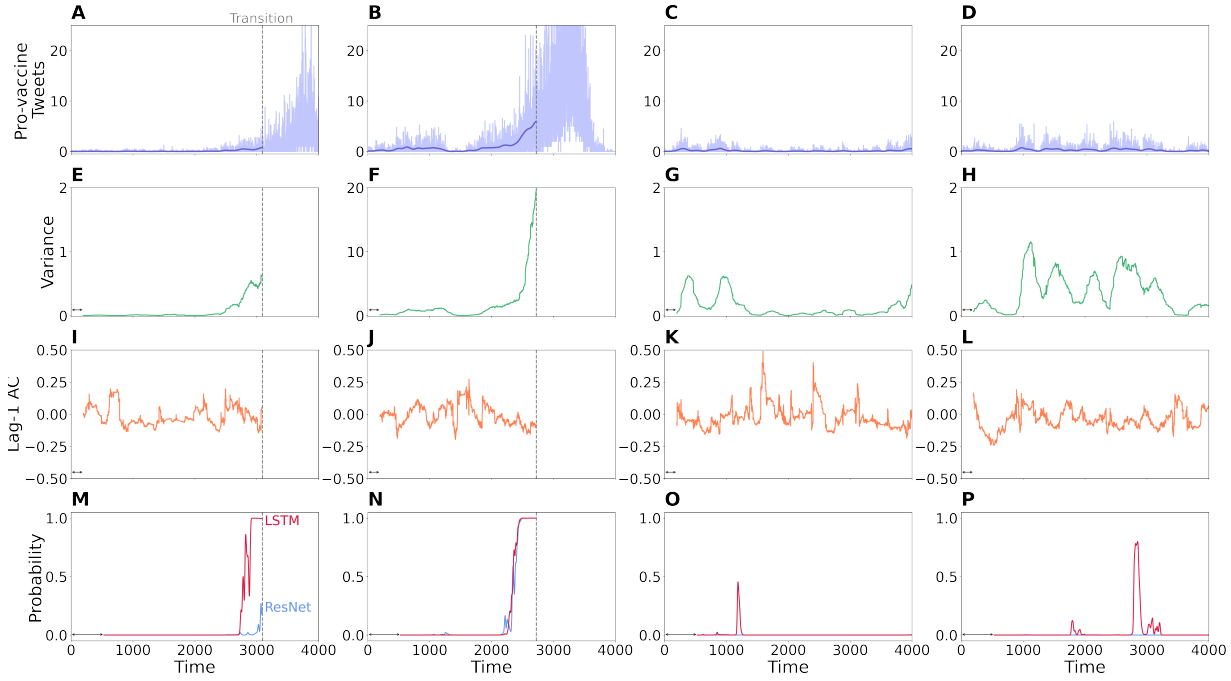


Figure 3.3: Trends in two conventional EWS (variance and lag-1 autocorrelation) and two deep learning classifiers (ResNet and LSTM) applied to simulated time-series data. (A–D) Corresponding number of pro-vaccine Tweets T_p (light purple) and its smoothed trend (dark purple) of Figures 3.2A–D. (E–H) Variance of T_p. (I–L) Lag-1 autocorrelation of T_p. (M–P) Probabilities assigned by two deep learning classifiers. In (A, B), grey dashed lines mark the transition time.

3.4.2 Performance of EWS and DL on empirical data

To test our deep learning classifiers on empirical data, we select four measles outbreaks in North America between 2013 and 2020 as positive samples: the 2014–15 outbreak in California, the 2018–19 outbreak in New York, the 2018–19 outbreak in Washington, and the 2014 outbreak in British Columbia, Canada (Figures 3.5A–D). For negative (or neutral) samples, we choose four countries (Singapore, Mexico, Argentina, and Brazil) that have mandatory measles vaccination policies and reported very few measles cases from 2014 to 2018 (Figures 3.6A–D).

Before the three outbreaks across the US (CA14, NY18, and WA19), both the ResNet and LSTM classifiers show a clear increase in probability predictions, signaling a potential

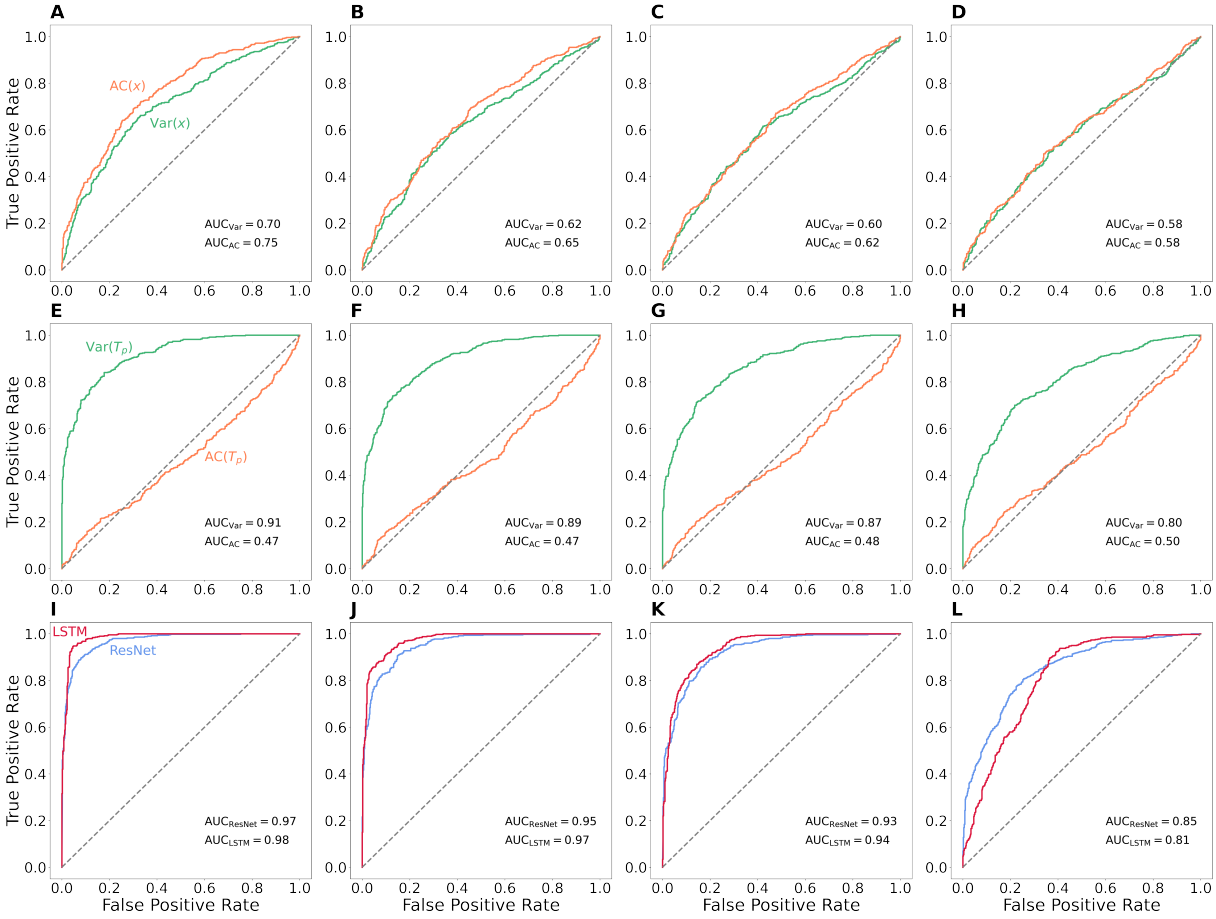


Figure 3.4: ROC curves for predictions of an upcoming critical transition in theoretical time-series data. (A–D) ROC curves based on Kendall tau values for variance and lag-1 autocorrelation of x. (E–H) ROC curves based on Kendall tau values for variance and lag-1 autocorrelation of T_p. (I–L) ROC curves for probabilities assigned by deep learning classifiers. Performance is evaluated on 500 positive and 500 neutral simulations. Panels correspond to different pre-transition periods used for prediction: (A, E, I) the entire pre-transition period, (B, F, J) up to 100 time points before the transition, (C, G, K) up to 200 time points before the transition, and (D, H, L) up to 400 time points before the transition. Grey dashed lines represent the chance level.

critical transition in the pro-vaccine percentage among the population (Figures 3.5M–O). Interestingly, while LSTM consistently responds faster than ResNet in the theoretical

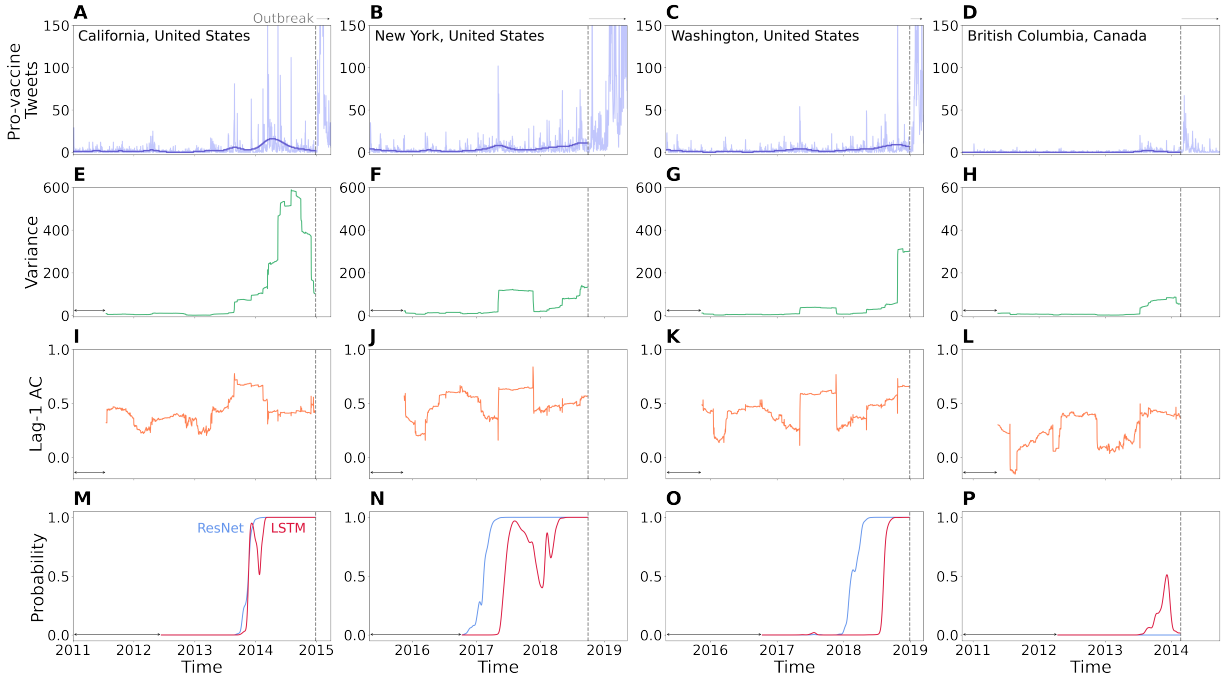


Figure 3.5: Trends in two conventional EWS and two deep learning classifiers applied to positive empirical Twitter data. (A–D) Number of pro-vaccine Tweets (light purple) and its smoothed trend (dark purple) before and during four selected outbreaks: (A) 2014–15 outbreak in California, (B) 2018–19 outbreak in New York, (C) 2018–19 outbreak in Washington, and (D) 2014 outbreak in British Columbia. Corresponding (E–H) variance, (I–L) lag-1 autocorrelation, and (M–P) probabilities assigned by two deep learning classifiers. Grey dashed lines indicate the start of each outbreak, with outbreak periods marked by arrows at the top of (A–D).

data, we find that ResNet is more sensitive to early signals in the empirical data. LSTM, however, proves effective in filtering out "false signals", such as sudden boosts in the number of Tweets caused by the reposting of news reports over a short period. In both CA14 and NY18 data, a small drop in the LSTM prediction occurs after it reaches near 100% (Figures 3.5M and N), as these short-lived bursts of discussion did not lead to a sustained discussion over time. Consequently, no significant drop in the current pro-vaccine percentage is expected. The number of Tweets collected in British Columbia is much smaller, making it challenging for deep learning classifiers to capture patterns in the time series data. The ResNet classifier predicts a zero probability for a critical transition

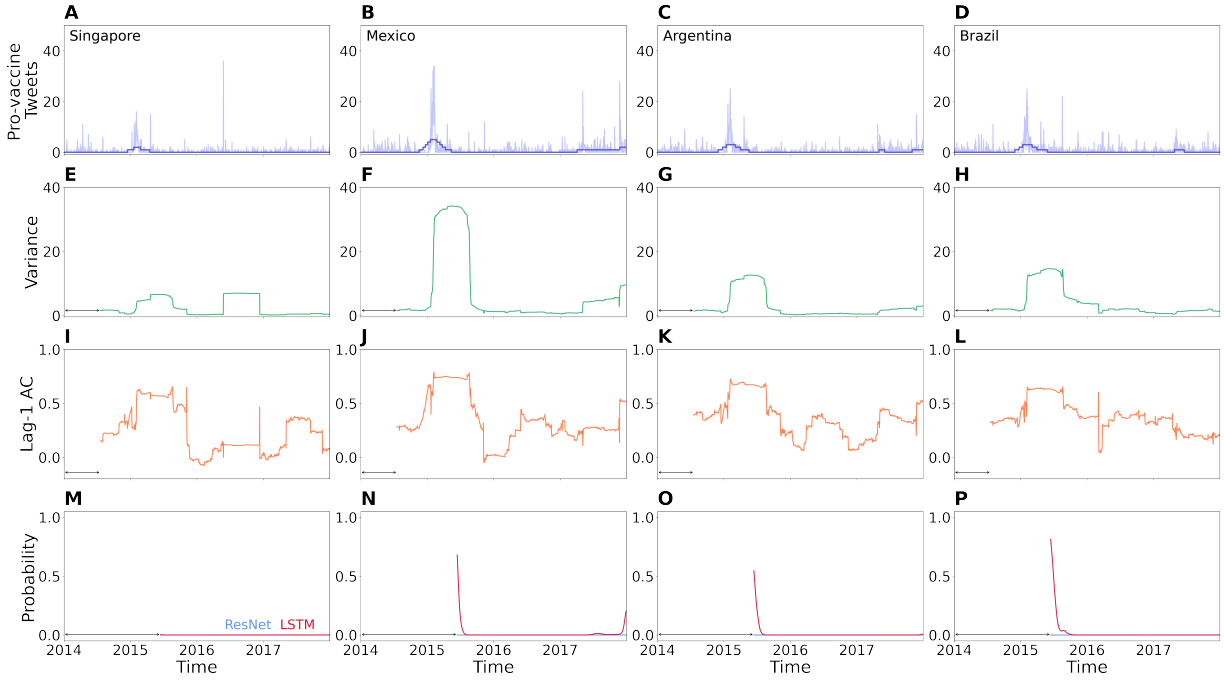


Figure 3.6: Trends in two conventional EWS and two deep learning classifiers applied to neutral empirical Twitter data. (A–D) Number of pro-vaccine Tweets (light purple) and its smoothed trend (dark purple) between 2014 and 2018 in four selected countries: (A) Singapore, (B) Mexico, (C) Argentina, and (D) Brazil. Corresponding (E–H) variance, (I–L) lag-1 autocorrelation, and (M–P) probabilities assigned by two deep learning classifiers.

to occur, while LSTM shows an increasing probability to around 55% during a period of moderate online discussion before dropping back to zero (Figure 3.5P).

In the four selected countries where measles vaccination is mandatory (SG, MX, AR, and BR), the number of pro-vaccine Tweets we collected is relatively small. Both ResNet and LSTM predict low probabilities of critical transitions throughout the time range (Figures 3.6M–P). In early 2015, all four countries experienced a relatively large-scale discussion following the outbreak at Disneyland, California. LSTM captures this slight increase in MX, AR, and BR data. However, the probabilities drop quickly as the discussion does not last long. Notably, no major measles outbreaks have occurred in these countries since then.

Our results reveal a marked increase in variance prior to each selected outbreak (Fig-

ures 3.5E–H). Lag-1 autocorrelation has an overall upward trend, with rises that start earlier but have a smaller scale (Figures 3.5I–L). For negative samples, both values exhibit only significant fluctuations when intense discussions are observed over a short period (early 2015); they remain relatively stable for the rest of the time (Figures 3.6E–L). In the case of variance, when we compare specific numerical values, these larger fluctuations are still much smaller compared to the values seen in positive cases.

3.5 Discussion

The coarse-scale nature of vaccination rate data poses challenges for public health departments in monitoring timely changes, thereby limiting their capacity to provide early warnings of potential declines. In contrast, social media provides high-frequency data on various topics, including vaccine-related sentiments. Our findings indicate that deep learning classifiers are adept at analyzing social media time series data to detect early warning signals of critical transitions in pro-vaccine sentiment. These classifiers not only exhibit superior sensitivity and specificity compared to conventional EWS like variance and lag-1 autocorrelation in theoretical data but also demonstrate robust performance when applied to empirical data, underscoring their practical utility in real-world scenarios.

Researchers have evaluated the performance of variance and lag-1 autocorrelation as EWS on a simple SEIRx (Susceptible-Exposed-Infectious-Recovered-Vaccinator) epidemic model [38]. In our study, we observe much lower AUC scores for both indicators on the pro-vaccine percentage, which we attribute to two key factors. First, unlike a model that broadly divides the entire population into pro-vaccine and anti-vaccine individuals, our model considers interactions within and between two distinct social groups. Additionally, the performance of two indicators is assessed not directly on a state variable but on a weighted average of two state variables with time-varying weights. These modifications introduce additional complexity to the system and can potentially distort the intrinsic statistical properties that conventional EWS rely on. Second, instead of Gaussian white noise, we apply fat-tailed α-stable Lévy noise that has an infinite variance when the stability parameter α is less than 2. This type of noise can produce large jumps that obscure the subtle and gradual changes typically expected near critical transitions, reducing the predictive power of variance. Furthermore, these large jumps can also disrupt the temporal relationship between a system's current state and its previous state, making lag-1 autocorrelation a less reliable indicator. The time series representing the number of pro-vaccine social media posts is more complex due to its derivation from transformed state variables through multiple functions, combined with additional Lévy noise. This complexity leads

to poor performance in lag-1 autocorrelation, which performs even worse than random guessing. On the other hand, the high performance of variance is expected, as the scale parameter of the additive Lévy noise is assumed to increase as the pro-vaccine percentage decreases [109], and such a decrease is a necessary condition for the system to approach a tipping point. Despite this complexity of data, our deep learning classifiers still serve as highly effective indicators. These findings highlight the advantage of using deep learning classifiers in capturing patterns in transformed and noisy data. This capability can be particularly valuable in real-world applications, as observational data often undergo multiple transformations and are influenced by non-Gaussian noise.

The performance of any deep learning model depends heavily on the quality of the training data [83]. For detecting early signals of pro-vaccine sentiment collapse on social media, the ideal training set should contain time series data from real-world social networks, paired with detailed vaccination data for the corresponding regions to determine whether a decline in vaccine coverage is occurring. However, collecting sufficient data from social media platforms is challenging, as most posts lack location information, and identifying vaccination rate declines from such coarse-scale data becomes even more difficult. Mathematical models provide an efficient way to simulate large amounts of training data; they also allow us to precisely control key conditions, such as whether a critical transition occurs and when it takes place. This approach has also been adopted in recent work, where theoretical epidemic models were used to generate synthetic data for training deep learning models to detect superspreading events [215]. Our approach uses a coupled behavior-disease model with two social groups as the backbone. We also introduce an exponential media function and randomize most of the parameters in order to make the data more realistic and diverse. However, it is important to acknowledge that our model relies on simplifications that may not fully capture the complexity of real-world social media behaviors. As a result, deep learning classifiers trained on those theoretical data could overfit to the synthetic patterns and struggle to generalize to the unpredictability of empirical data. To address this sim-to-real gap, a potential direction is to pre-train the model on simulated data, followed by fine-tuning on limited real-world data (a specific form of transfer learning) [122]. This procedure usually allows models to retain structural insights from simulation while adapting to empirical noise.

The complex dynamics of social media data require sufficient temporal context to model effectively. As a result, our approach is particularly well-suited for regions that have not experienced outbreaks for extended periods, during which vaccination rates are in decline. Ideally, testing samples should come from regions with outbreaks in the early 2010s, when the resurgence of measles began [84, 29]. However, since social media did not gain widespread popularity until the 2010s, few posts were collected prior to these

outbreaks, making the time series data less informative. For outbreaks occurring in the mid-to-late 2010s, the preceding social media data was often influenced by outbreaks in other regions. To address this, we reduce the length of the time series to isolate a cleaner segment of data, minimizing the impact of sudden boosts caused by outbreaks elsewhere. Our choice of 500 data points, corresponding to approximately 1.5 years of daily data, strikes a balance between capturing annual trends and seasonal effects while maintaining computational feasibility. However, a longer time series would be preferable for detecting EWS [30]. To achieve a more accurate reflection of vaccination rate fluctuations, more detailed data preprocessing methods will be required to filter out the effects of outbreaks on social media data.

In the four countries with mandatory measles vaccination policies included in our study, the number of Tweets related to the measles vaccine was low for two main reasons. First, the policy has made vaccination a socially accepted norm, resulting in little public debate. Second, a large portion of the discussions occurred in languages other than English, since English is not an official or commonly spoken language in Mexico, Argentina, and Brazil, although it is one of Singapore's four official languages. Therefore, the English Tweets we collected represent only a small fraction of the total discussions in these countries, which introduces a selection bias and leads to underrepresentation. Future research could adopt multilingual approaches using language-specific sentiment models. This would also help expand the real-world dataset, which could support more effective transfer learning through pre-training and fine-tuning.

We proposed two widely used deep learning frameworks for time series classification. While each has its strengths, we do not consider one approach strictly better than the other. The LSTM model is more computationally efficient due to a much smaller number of parameters. It generally outperforms ResNet but tends to overreact to short-term fluctuations in some cases. The ResNet model is more conservative and predicts a high probability of critical transitions only when enough evidence is observed. We believe that combining insights from both models offers a more balanced and comprehensive perspective, leveraging the sensitivity of LSTM and the stability of ResNet for more robust analysis of time series data.

We hope this study provides new insights into predicting critical transitions in dynamical systems, and contributes to the growing field of behavioral epidemiology, particularly in the context of parental vaccination choices. Human behaviors introduce significant complexity and noise into these systems, often limiting the effectiveness of traditional statistical early warning indicators. Additionally, real-world data shortages can also make theoretically feasible models difficult to implement in practice. Our findings demonstrate that it is possible to bridge the gap between dynamical systems and data-rich sources like social

media and let deep learning models extract meaningful patterns from available data and uncover early signals of critical transitions, even in noisy and complex environments. Integrating these models with real-time social media monitoring could serve as a powerful tool for proactive public health interventions, ultimately helping to reduce the risks of disease outbreaks.

Chapter 4

Homophily in a multi-group vaccination game and its impact on disease dynamics

This chapter is based on the paper: Z. He, P. Manfredi, A. d'Onofrio, and C. T. Bauch. Homophily in a multi-group vaccination game and its impact on disease dynamics. *Proceedings of the Royal Society B: Biological Sciences*, 2025. Under review.

4.1 Abstract

Metapopulation models are among the most frequently used approaches to address spatial heterogeneity in infectious disease modeling. While researchers typically consider the role of heterogeneity in disease transmission by dividing the population into subgroups according to levels of social activity, susceptibility or perceived risk of infection, the role of heterogeneity in vaccination behavior has received less attention. The tendency to interact with others who share similar beliefs (homophily) can further reinforce such heterogeneity, leading to clustered vaccine sentiments in the population. In this paper, we develop a multi-group vaccination game model coupled with an SIRV epidemic model and examine how interactions among individuals in–and between–groups affect both vaccination behavior and disease dynamics. We find that in a three-group system with extremely pro- and anti-vaccine groups, each group can converge to a multi-strategy equilibrium, causing population-level opinion polarization that can persist despite the severity of the disease. Our results highlight the dual role of homophily, offer insights on the design of effective public health interventions, and demonstrate the potential of a multi-group framework in real-world disease modeling.

4.2 Introduction

Vaccination has been one of the most effective ways to control the spread of infectious diseases, as it reduces the size of outbreaks by providing immunity to individuals. Before the measles vaccine was introduced in the United Kingdom (UK) in 1968, major outbreaks occurred in the UK every one to two years in a non-periodic fashion, causing hundreds of deaths annually [220, 77]. More complex patterns were observed in some other countries [217, 129, 34]. With the help of mass vaccination programs, measles epidemics became less frequent and smaller in scale. However, in 1998, a series of fraudulent papers sparked widespread fear about a link between the MMR (Measles, Mumps, and Rubella) vaccine and autism [125, 35]. This led to a decline in vaccine uptake and the resurgence of measles in the UK, highlighting how the success of vaccination programs greatly depends on people's willingness to get vaccinated.

Vaccination is not just a pharmaceutical intervention but also a strategic decision. People have to weigh the risk of side effects from vaccines against the often larger risk of infection, which can cause severe illness and long-term health complications. As a result, vaccination is often a more attractive strategy during a pandemic, but when a disease becomes rare, choosing not to vaccinate may seem more appealing. Herd immunity (HI)

further complicates this decision. A widely diffused misinterpretation of HI is thinking that it eliminates all risk of infection, which is incorrect. In fact, HI only means that the disease cannot persist in a highly immunized population, although localized, sporadic infections can still occur. Since a disease can be eliminated without vaccinating everyone, herd immunity creates opportunities for free riders by allowing some people to avoid vaccination while still benefiting from protection. As more people are vaccinated or naturally immunized, the incentive for others to get vaccinated decreases. However, if too many people take this approach, the overall immunity of the population can drop below the level needed to prevent outbreaks.

Integrating vaccination behavior into epidemic models has become a common approach for studying the spread of vaccine-preventable diseases [93, 152, 229]. These models usually include two interconnected layers of dynamics: a biological layer that describes how the disease transmits in a population and a social layer that captures how individuals imitate others and adjust their vaccination strategies [228]. The biological layer is typically represented using compartmental models, such as the classic SIR (Susceptible-Infected-Recovered) framework or its variations. In these models, individuals are grouped into compartments based on their infection status, and differential equations describe the rates at which individuals move between compartments [7]. The social layer, on the other hand, can be modeled using evolutionary game theory [85, 223, 48]. It helps explain vaccination behavior by treating vaccination as a "game", where individuals make decisions based on both self-interests as well as the choices of others. In the meantime, it captures how vaccination decisions evolve over time, as people continually adjust their strategies according to the success or failure of other people's strategies.

A substantial body of research explores disease-behavior dynamics in populations where homogeneous mixing is assumed [24, 22, 20, 28, 21, 72, 173]. This assumption simplifies the model structure, allowing researchers to easily analyze key outcomes, such as equilibrium points, stability conditions, and threshold parameters like the basic reproduction number R_0. These models also give insights into the fundamental interplay between behavior and disease dynamics before introducing more complex factors. However, such an assumption overlooks important features of real-world populations, such as geographic and social structures that could hugely affect the spread of the disease and opinions [17, 169, 218]. To consider heterogeneity in populations, researchers often apply complex networks, where individuals are represented as nodes in a network, with edges in different layers denoting interactions that enable disease transmission as well as social contacts [228, 169, 104, 9, 227, 245, 137].

Patch models provide another useful way to capture population heterogeneity by dividing individuals into distinct subgroups [189]. Unlike complex networks that empha-

size individual-level contact patterns to trace transmission, patch models focus on group-level interactions to capture dynamics between subpopulations, making them easier to analyze, particularly when detailed behavioral data are unavailable, which is the most common scenario. A normal practice is to group people by age to reflect differences in social contact patterns, susceptibility to infection, and vaccine responses, which has been widely used in designing targeted vaccination strategies for seasonal influenza and COVID-19 [120, 143, 243]. Researchers have also assigned different levels of risk perception [55, 54, 18, 205] and vaccine efficacy to subgroups [50, 102] to account for diverse characteristics of the population. However, existing models typically consider heterogeneity only on the biological layer, while its effect on the social layer is often ignored. In reality, vaccine sentiments tend to cluster in social and geographic space, creating communities with high or low vaccine acceptance [194, 127]. This suggests that a single variable cannot adequately capture the population-wide pro-vaccine sentiment. Moreover, with the widespread use of social media and online messaging tools, people have more channels to access information, and the speed at which they acquire information has greatly increased. People who spend more time online may be more exposed to vaccine-related information, whether accurate or misleading [11, 148, 13], while those with limited or no Internet use may rely more on traditional social contacts. As a result, imitation dynamics in vaccination behavior can vary significantly across groups. These observations highlight the need to account for heterogeneity at the group level in models of vaccination behavior.

In particular, homophily (or assortativeness, in the language of epidemiological modeling), the tendency for individuals to socially interact with those who share similar views or behaviors [157, 88, 7, 166], can strongly influence how vaccination strategies or other health behaviors spread in the population [193, 47, 211]. A previous study divides the population into social media users and non-users and explores how homophily and online events affect vaccination behavior as well as their effect on disease transmission [109]. In this study, we extend the model to more than two groups. Our aim is to examine how small, highly committed pro- and anti-vaccine minority groups interact with the broader majority of the population. Specifically, we ask whether the majority acts as a swing group that shifts under the influence of these extreme groups, or whether the minority groups gradually conform to the prevailing views of the majority over time.

4.3 Methods

We propose a coupled behavior-disease model for a pediatric infection such as measles, where vaccination is typically administered at birth [20, 72]. The disease (or biological)

layer follows the Susceptible-Infected-Recovered-Vaccinated (SIRV) framework with natural births and deaths, under the assumption of a homogeneous population. However, on the social layer, we divide the population into n distinct groups of various sizes but assume well-mixing within each subgroup. In the post-COVID-19 era, changes in lifestyle such as reduced travel and increased remote working have altered the dynamics of both disease spread and information flow. As a result, even when considering a small, homogeneous community from a disease transmission perspective, social heterogeneity remains significant. While heterogeneity on the biological layer may be negligible in such cases, differences in information dissemination across social groups must be considered.

On the biological layer, individuals are classified as susceptible (S), infected (I), recovered (R), or vaccinated (V) based on their health status. On the social layer, parents of newborn children act as decision-makers in the "vaccination game" [22, 20]. Each parent belongs to one of the n social groups and is either pro-vaccine (P) or anti-vaccine (A). We assume that parents strictly adhere to their vaccination strategy and that vaccinated children would gain lifelong immunity. At each time point, the proportion of vaccinated newborns matches the proportion of pro-vaccine parents in the entire population (i.e., the weighted average across all social groups). Effectively vaccinated children enter the system as vaccinated individuals, whereas the rest enter as susceptible.

We consider the presence of injunctive social norms, where individuals are rewarded for aligning with the majority within their group [112, 113, 173]. However, our primary focus is on homophily. We assume that homophily levels vary across groups and that homophily influences imitation dynamics in two key ways. First, social norms from other groups have less influence. Second, imitation between individuals from different groups occurs less frequently [109].

We let x_i be the proportion of pro-vaccine parents in group i, and thus the proportion of anti-vaccine parents in the same group is $1 - x_i$. The payoffs for adopting a pro-vaccine and an anti-vaccine strategy in group i are

$$\pi_i(P) = -r_{\mathrm{v}} + \delta_0 \Big[x_i + (1 - h_i) x_{-i} \Big],$$
$$\pi_i(A) = -cI(t) + \delta_0 \Big[1 - x_i + (1 - h_i)(1 - x_{-i}) \Big]$$

with

$$x_{-i} = \sum_{j \neq i} \frac{N_j x_j}{N - N_i}$$

indicating the external pro-vaccine percentage for group i. r_{v} is the perceived risk from the vaccine and the term $cI(t)$ represents the perceived risk of infection, where c is a

proportionality constant and the risk is assumed to increase linearly with the proportion of infected individuals $I(t)$. δ_0 quantifies the effect of injunctive social norms, which is assumed to be the same in all groups. $h_i \in [0,1]$ represents the level of homophily in each group, where higher values indicate stronger homophily, meaning individuals interact less frequently with people from different groups. Said otherwise, the level of homophily is a measure reflecting the preference for socially interacting with individuals of their own group. This has motivated the key concept of preferred mixing in epidemiological modeling [115].

Given that the principal focus of this work is on the "pure" effects of homophily, our multi-group model disregards any other possible sources of heterogeneity. For this reason, we assume that all other parameters related to the diffusion of opinions on vaccination, namely perceived costs arising from infection or vaccination as well as the strength of social norms, are constant across all social groups. Though easy to include, these additional heterogeneities would make the overall picture less sharp.

We adopt important assumptions consistent with previous work that individuals switch strategies only when doing so yields a higher payoff, and the imitation rate is proportional to the resulting payoff gain. In other words, individuals are more likely to switch strategies if the potential payoff improvement is larger. Given the social learning rate κ, the rate of change in the pro-vaccine proportion x_i, accounting for imitation between group i and group j, is [111]

$$\kappa x_j(1-x_i)\max[\pi_j(P)-\pi_i(A),0]-\kappa(1-x_j)x_i\max[\pi_j(A)-\pi_i(P),0],$$

where $\max[z,0]$ denotes the positive part of z, also known as the ReLU function. When $j=i$, this expression simplifies to the within-group imitation rate:

$$\kappa x_i(1-x_i)[\pi_i(P)-\pi_i(A)].$$

We simplify the notation by rescaling r_{v}, c, δ_0, and κ such that $\omega \equiv r_{\mathrm{v}}/c$, $\delta \equiv \delta_0/c$, and $\kappa \equiv \kappa c$. For clarity, we use ω, δ, and κ throughout the paper.

To incorporate homophily, we make two additional assumptions: (i) the total amount of learning is fixed, and (ii) increasing the number of groups does not increase the total amount of learning but instead distributes it across the groups. To reflect this, we assign a weight of $[1+(n-1)h_i]/n$ for imitation within the group and $(1-h_i)/n$ for imitation between groups. This weighting ensures that the total amount of learning remains constant:

$$\left(\frac{1+(n-1)h_i}{n}\right)\kappa+(n-1)\left(\frac{1-h_i}{n}\right)\kappa=\kappa.$$

Combining these elements, we obtain the system of ordinary differential equations describing the coupled behavior-disease dynamics:

$$\frac{dS}{dt} = \mu(1-\epsilon x) - \beta SI - \mu S, \tag{4.1}$$

$$\frac{dI}{dt} = \beta SI - \gamma I - \mu I, \tag{4.2}$$

$$\frac{dR}{dt} = \gamma I - \mu R, \tag{4.3}$$

$$\frac{dV}{dt} = \mu\epsilon x - \mu V, \tag{4.4}$$

$$\begin{aligned}\frac{dx_i}{dt} = &\left(\frac{1+(n-1)h_i}{n}\right)\frac{\kappa N_i}{N}\Big\{x_i(1-x_i)[\pi_i(P)-\pi_i(A)]\Big\} \\ &+\left(\frac{1-h_i}{n}\right)\sum_{j\neq i}\frac{\kappa N_j}{N}\Big\{x_j(1-x_i)\max\left[\pi_j(P)-\pi_i(A),0\right] \\ &-(1-x_j)x_i\max\left[\pi_j(A)-\pi_i(P),0\right]\Big\}.\end{aligned} \tag{4.5}$$

Given the population sizes N_i, the overall proportion of pro-vaccine individuals (parents) in the population is calculated as

$$x = \frac{\sum_i x_i N_i}{N}.$$

Here, κ also accounts for the interaction rate, and the factor N_i/N ensures proportionality to the relative frequency of group i [111]. Since $S + I + R + V = 1$ always holds, the differential equation for R can be eliminated. The parameters used in this system are summarized in Table 4.1.

Throughout this study, we assume that the population mainly consists of a large "general group". The remaining individuals who hold strong pro-vaccine or anti-vaccine sentiments are divided into several "extreme groups". To begin, we consider a population with three subgroups ($n = 3$): the general group ($i = 1$), one pro-vaccine group ($i = 2$), and one anti-vaccine group ($i = 3$). The two minority groups ($i \in \{2, 3\}$) are assumed to exhibit relatively high homophily, reflecting the empirical fact that individuals preferentially interact with others who share similar views, especially in the vaccine debate [163]. More details on numerical methods are provided in Appendix B.3.

A comprehensive model in evolutionary game theory accounts not only for imitation processes but also for other types of transitions [111]. That is, individuals can switch strategies independently of social learning. For instance, vaccination campaigns led by

Parameter	Description	Baseline value	References
n	Number of social groups	3	Assumption
N	Total population size	100000	Assumption
N_i	Population size of subgroups	$[80000, 10000, 10000]$	Assumption
R_0	Basic reproduction number	10	[20, 105]
γ	Recovery rate	24/year	[7, 77, 20]
μ	Natural birth/death rate	1/60/year	[7, 77, 20]
β	Transmission rate	240/year	[20, 173]
ϵ	Vaccine efficacy	0.9	[21]
κ	Social learning rate	1000	[175]
ω	Relative vaccine risk	0.0003	[175]
δ	Injunctive social norm strength	0.0008	[173, 175]
h_i	Homophily level	$[0.5, 0.8, 0.8]$	Assumption

Table 4.1: Parameters and descriptions. The transmission rate β is calculated by $R_0(\gamma+\mu)$.

the public health system (PHS) can directly influence anti-vaccine individuals to adopt pro-vaccine behaviors [72]. In addition, individuals can switch between strategies based on purely intrinsic motivations. Accordingly, we propose a more general framework that extends Eq. (4.5) to incorporate these types of transitions:

$$\begin{aligned}\frac{dx_i}{dt} = &\left(\frac{1+(n-1)h_i}{n}\right)\frac{\kappa N_i}{N}\Big\{x_i(1-x_i)[\pi_i(P)-\pi_i(A)]\Big\} \\ &+\left(\frac{1-h_i}{n}\right)\sum_{j\neq i}\frac{\kappa N_j}{N}\Big\{x_j(1-x_i)\max\left[\pi_j(P)-\pi_i(A),0\right] \\ &-(1-x_j)x_i\max\left[\pi_j(A)-\pi_i(P),0\right]\Big\} \\ &+\lambda_i(t)(1-x_i)-\epsilon_i x_i, \end{aligned} \tag{4.6}$$

where $\lambda_i(t)$ is a constant or a periodic positive function, and ϵ_i is a constant. When λ_i is a small constant such that $\lambda_i = \epsilon_i$, we face a situation of random exploration of the other strategy by individuals in each opinion group. Notably, Eq. (4.5) can be obtained by setting $\lambda_i = \epsilon_i = 0$ in Eq. (4.6). A detailed derivation can be found in Appendix B.1.

4.4 Results

4.4.1 Behavioral dynamics in the absence of diseases

Although global measles eradication has not been achieved and elimination has only been sustained at the national level, in some smaller areas, such as certain cities and counties, measles cases are extremely rare and can be considered eliminated. Therefore, it is valuable to consider a baseline scenario in which the disease is absent (i.e., $I(0) = 0$), implying that $I(t) = 0$ for all future t. In this case, vaccination dynamics are driven purely by behavioral and social influences.

When a single well-mixed group is assumed on the social layer, the disease-free equilibrium with partial vaccine coverage $x^* = (1 + \omega/\delta)/2$ is always unstable whenever it exists, since the social norm drives the proportion of pro-vaccine individuals either up to 1 or down to 0 [173]. Similarly, if we consider multiple groups but assume full homophily within each group, meaning individuals only imitate members of their own group, the dynamics within each group become independent. Each group behaves like an isolated well-mixed population, and the instability of partial vaccine coverage persists within each group.

However, when the population is divided into multiple groups and cross-group imitation is allowed, the dynamics become more complex, and the previous threshold condition may no longer hold. If the pro-vaccine percentage in the out-group is lower than that in the in-group, excessive cross-group interaction may actually hinder vaccine uptake, raising the critical threshold above $(1 + \omega/\delta)/2$. In contrast, when the out-group has a very high pro-vaccine percentage, even if the initial in-group level $x_i < (1 + \omega/\delta)/2$, cross-group imitation may still drive x_i up to 1. On top of this, homophily introduces an additional layer of complexity, as it can amplify or dampen the influence of other groups' behaviors.

Figure 4.1 illustrates the interplay between initial pro-vaccine percentages and homophily of the majority. We consider two initial pro-vaccine levels that fall below the critical threshold in a scenario of full homophily in the general group. When the opinion flow within the group tends to shift towards anti-vaccine strategies, low homophily, allowing frequent interactions between groups, can temporarily stabilize or even reverse the decline (Figure 4.1*a* and *e*). When the initial pro-vaccine level exceeds the critical threshold and the group favors pro-vaccine strategies, low homophily tends to accelerate the approach to full vaccination (Figure 4.1*i*).

The behavior of the general group also significantly influences the dynamics of the other two minority groups, despite the high homophily levels, since most of the interactions oc-

cur with general group members (Figure 4.1*b*, *c*, *f*, *g*, and *j*). The trend in the effectively vaccinated proportion is largely driven by the behavior of the general group as well. However, due to the relatively slow rates of natural birth and death, changes in vaccination coverage occur gradually rather than abruptly (Figure 4.1*d*, *h*, and *l*). We observe that in a high homophily setting, the increase in x_1 fails to drive up the pro-vaccine percentage in the highly anti-vaccine group (Figure 4.1*k*). This is because the high level of homophily reduces the payoff of vaccination, making it a less appealing strategy.

To further investigate how group-specific homophily levels shape long-term vaccination outcomes, we analyze the joint effect of the homophily in the two minority groups on the final proportion of pro-vaccine individuals (Figure 4.2). We select two initial conditions for the general group's pro-vaccine percentage, 0.67 and 0.69, lying on either side of the critical threshold $(1+\omega/\delta)/2$. Lower initial values of x_1 are avoided, as they tend to drive all groups towards a fully anti-vaccine state across almost all combinations of h_2 and h_3 within their targeted range, except when h_2 is very small or becomes unrealistically high. From Figure 4.2*a* and *b*, the system consistently converges to a fully pro-vaccine state across all groups when both h_2 and h_3 are small. As h_3 increases, the highly anti-vaccine group becomes increasingly entrenched in vaccine refusal, often resulting in a multi-strategy steady state in the population. A rise in h_2 reinforces the pro-vaccine dominance within the group at the beginning, but its long-term trend depends on the general group.

In addition to the influence of group-specific homophily, Figure 4.2*a* and *b* also reveal a pronounced sensitivity to initial conditions near the critical threshold: small differences in the initial pro-vaccine percentage in the general group can lead to dramatically different long-term outcomes. When $x_1(0)$ increases from 0.67 to 0.69, the region in which the system converges to a fully pro-vaccine equilibrium (yellow) expands, indicating that the requirements to achieve widespread vaccine uptake become less strict. The system can tolerate higher homophily levels in both minority groups while still reaching a favorable outcome. The time series in Figure 4.2*c*–*h* illustrate how such a modest shift in initial conditions can qualitatively alter the equilibrium of the system. In some cases, the system changes from a fully anti-vaccine state to a partially pro-vaccine one (Figure 4.2*c* and *d*), or from a partially pro-vaccine state to a fully pro-vaccine one (Figure 4.2*e* and *f*). A more extreme case where the system directly shifts from a fully anti-vaccine equilibrium to a fully pro-vaccine one is also possible (Figure 4.2*g* and *h*). Notably, in Figure 4.2*f* and *h*, the time series for x_3 are not monotone. They initially decrease before eventually increasing, indicating the possibility of transient declines in vaccine uptake before it recovers.

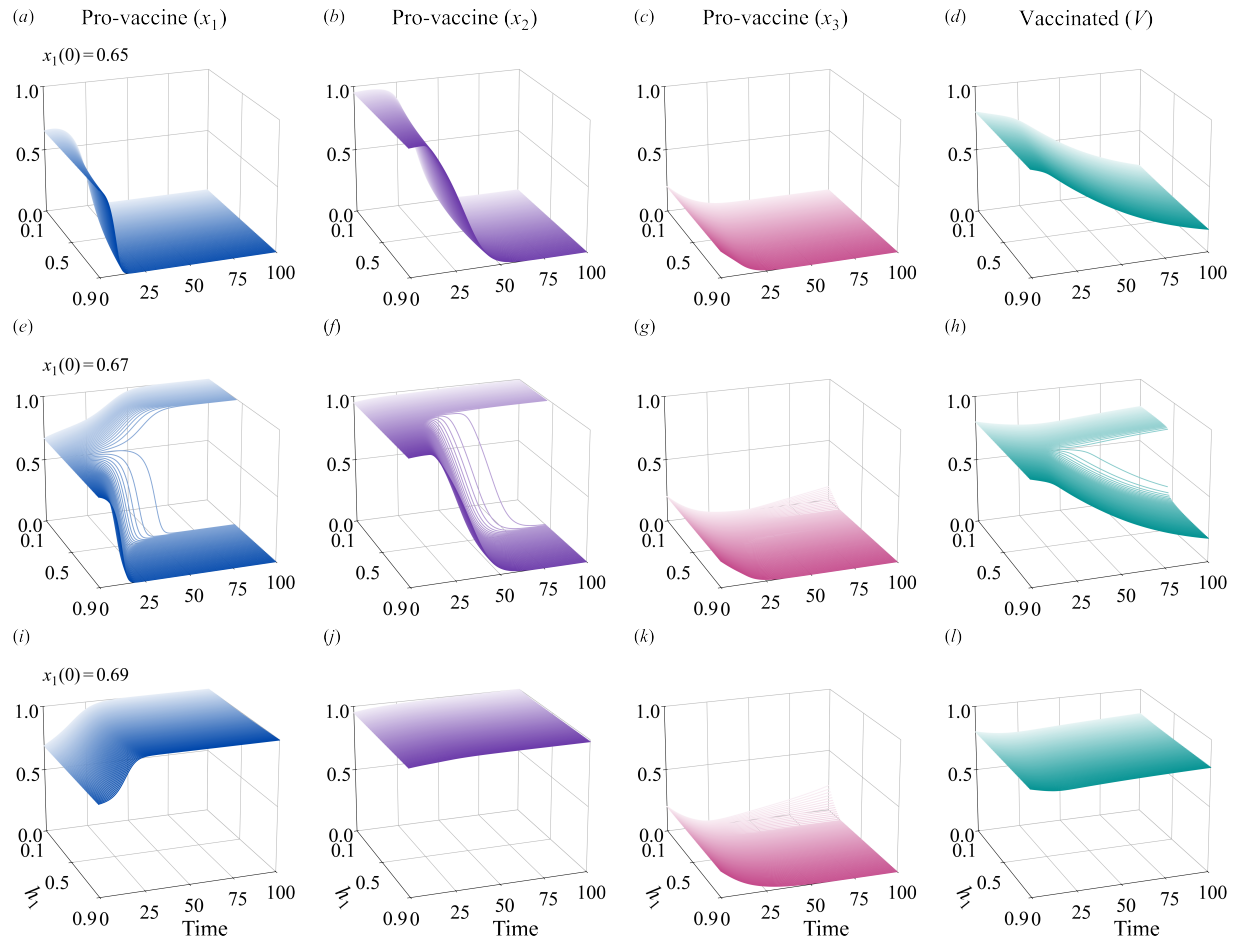


Figure 4.1: Dynamics of pro-vaccine opinions and vaccine coverage under different levels of homophily in the general group when the disease is absent. Each panel shows the temporal evolution of the proportion of pro-vaccine individuals in (a, e, i) the general group, (b, f, j) the highly pro-vaccine group, and (c, g, k) the highly anti-vaccine group, as well as (d, h, l) the overall vaccinated percentage of the entire population. The rows correspond to different initial pro-vaccine percentages in the general group: (a–d) $x_1(0) = 0.65$, (e–h) $x_1(0) = 0.67$, and (i–l) $x_1(0) = 0.69$.

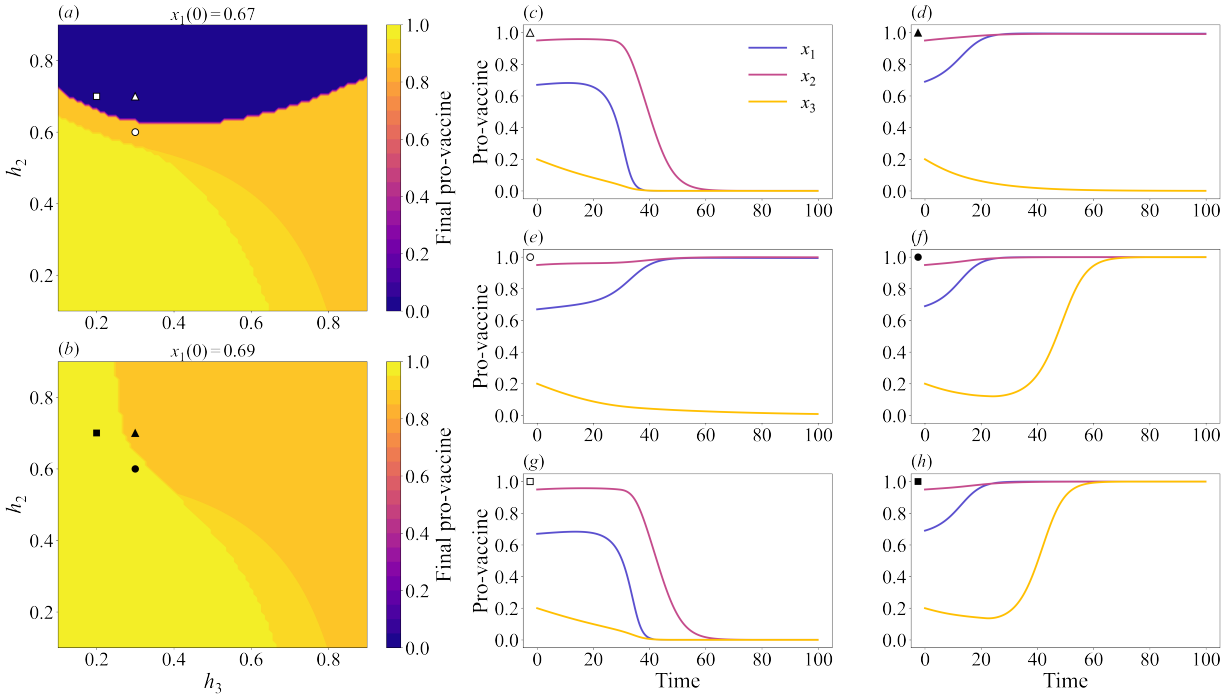


Figure 4.2: Impact of homophily levels in the minority groups on final pro-vaccine proportion in the entire population when the disease is absent. (a, b) Final pro-vaccine proportion across combinations of homophily levels, with initial pro-vaccine proportion in the general group set to $x_1(0) = 0.67$ and $x_1(0) = 0.69$, respectively. (c–h) Time series plots of pro-vaccine percentages in three groups, corresponding to (c, e, g) the white markers in (a), and (d, f, h) the black markers in (b). The homophily combinations illustrated are: $h = [0.5, 0.7, 0.3]$ (triangle), $h = [0.5, 0.6, 0.3]$ (circle), and $h = [0.5, 0.7, 0.2]$ (square).

4.4.2 Impact of time-varying homophily levels under endemic conditions

In real-world social systems, homophily is hardly a static characteristic. The strength of social ties within and between groups can evolve in response to external factors such as media influence and public health interventions, especially when the disease is present. Social media platforms and search engines often reinforce echo chambers by showing users content similar to their existing views [53], and this feedback loop strengthens group identities and isolates users from opposing perspectives. On top of this, individuals are more likely to seek social interactions that confirm their beliefs, known as the confirmation bias [160, 151].

These phenomena typically cause homophily to become stronger over time. On the other hand, homophily can also weaken under certain conditions, such as through educational campaigns by public health agencies and algorithmic adjustments of social media platforms that promote exposure to diverse content.

To explore how such temporal changes in social connectivity affect coupled behavior-disease dynamics, we compare, under endemic conditions, scenarios where the homophily level h_i increases or decreases over the medium/long term against a baseline scenario where homophily remains constant over time. For increasing homophily, we use the logistic function

$$h_i(t) = \frac{2(L - h_i^{\mathrm{b}})}{1 + e^{-\alpha t}} + 2h_i^{\mathrm{b}} - L,$$

where h_i^{b} is the baseline value, $L = 0.9$ is the upper bound, and $\alpha > 0$ controls the rate of increase. Similarly, for decreasing homophily, we use

$$h_i(t) = \frac{2(h_i^{\mathrm{b}} - L)}{1 + e^{\alpha t}} + L,$$

where $L = 0.1$ now serves as the lower bound and α controls the rate of decrease. In the two scenarios with evolving homophily, we consider two distinct values of the time scale of homophily evolution, that is, 20 and 50 years.

We begin by examining the baseline case of constant homophily ($\alpha = 0$), shown in Figure 4.3*a*–*c* (and also in Figure 4.4*a*–*c*, where the same baseline case is included for comparison). For simplicity, we refer only to Figure 4.3 in the following discussion. The dynamics observed in the presence of disease (Figure 4.3) share important similarities with the disease-free scenario (Figure 4.2), where the population can reach an anti-vaccine consensus (Figure 4.3*a*) or become more polarized (Figure 4.3*b* and *c*) in the long run. However, the trajectories in Figure 4.3 show oscillations, corresponding to repeated disease outbreaks that cause sudden boosts in pro-vaccine sentiments. Another notable difference is that, when $x_1(0) = 0.67$, the presence of disease fundamentally alters the system's outcome. In the endemic setting, outbreaks drive x_1 and x_2 up, eventually pushing them toward 1 (Figure 4.3*b*); when the disease is absent, the same initial condition would lead to declines in both variables, with both approaching 0.

Our results show that as homophily strengthens, the influence of the general group through high contact rates on minority groups diminishes. In the anti-vaccine group, support for vaccination remains very low, even if the support rate in the general group eventually approaches 1. In contrast, maintaining a high level of vaccine support within the pro-vaccine group becomes easier. When the support rate in the general group decreases

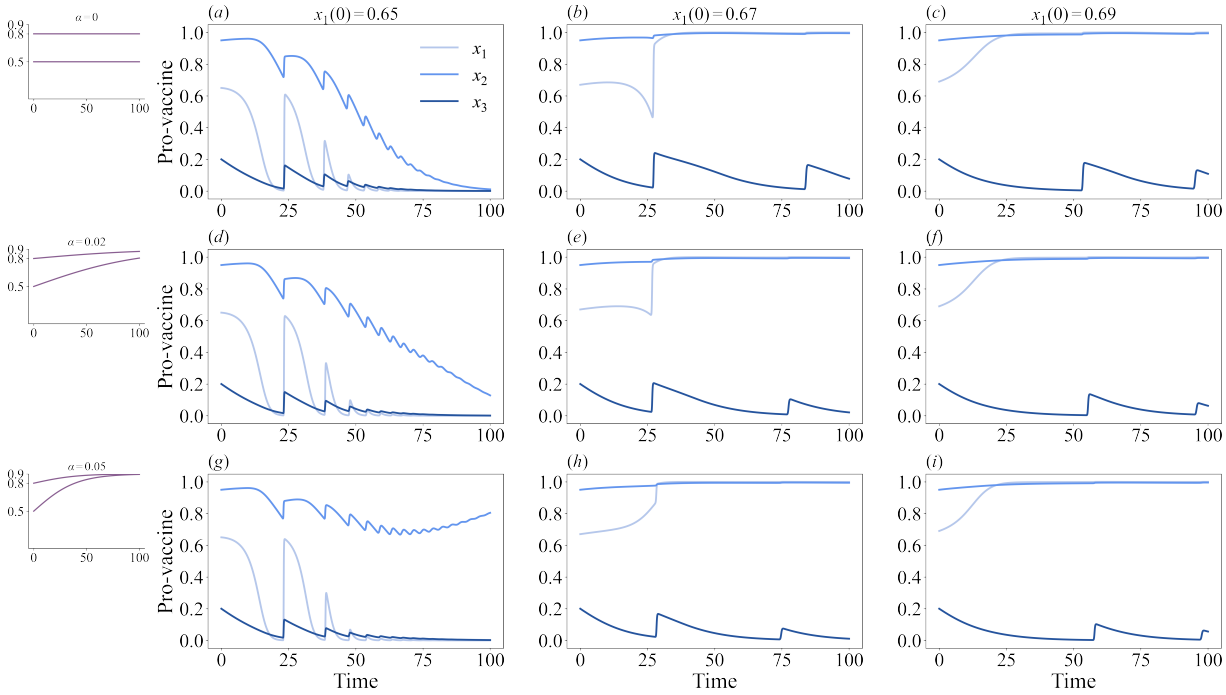


Figure 4.3: Dynamics of pro-vaccine opinions across three groups under constant and logistically increasing homophily levels in the presence of the disease. Each row corresponds to a different rate of decline: $(a$–$c)$ $\alpha = 0$, $(d$–$f)$ $\alpha = 0.02$/year, and $(g$–$i)$ $\alpha = 0.05$/year. Each column corresponds to a different initial opinion level in the general group: $(a,\ d,\ g)$ $x_1(0) = 0.65$, $(b,\ e,\ h)$ $x_1(0) = 0.67$, and $(c,\ f,\ i)$ $x_1(0) = 0.69$. The small panels on the left show the time-varying homophily levels for each α. Oscillations in the trajectories represent epidemics (for which graphs are not reported), leading to fast opinion switches over a rather short period of time in each social group.

rapidly and fluctuates at low levels, the support rate in the pro-vaccine group can stabilize or even rebound, increasing back to 1 (Figure 4.3).

When homophily levels decline, we observe that the entire population eventually reaches consensus, becoming either fully pro-vaccine or fully anti-vaccine, and the time it takes to reach this consensus decreases as the rate of decline increases. As the difference between within-group and cross-group influence becomes smaller, the high contact rate enables the general group to dominate the dynamics of minority groups. Specifically, vaccine support in the anti-vaccine group can no longer remain at a low level independently of other groups, while the support in the pro-vaccine group, once substantially reduced, does not bounce

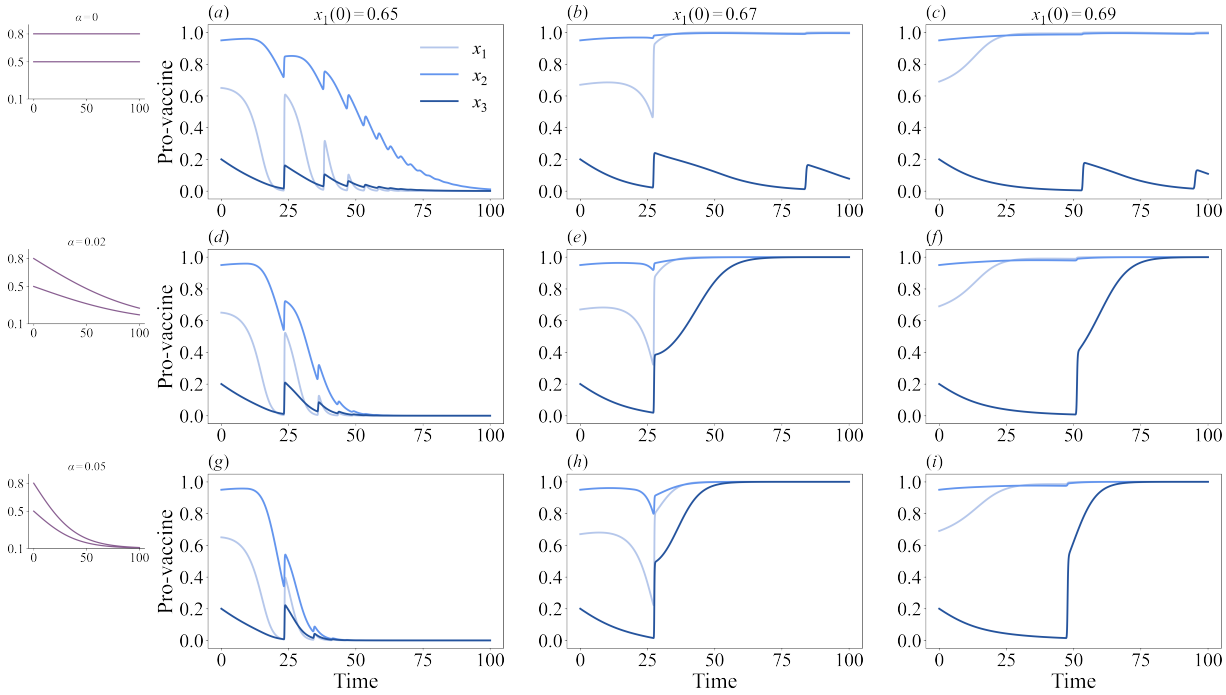


Figure 4.4: Dynamics of pro-vaccine opinions across three groups under constant and logistically decreasing homophily levels in the presence of the disease. Each row corresponds to a different rate of decline: $(a$–$c)$ $\alpha = 0$, $(d$–$f)$ $\alpha = 0.02$/year, and $(g$–$i)$ $\alpha = 0.05$/year. Each column corresponds to a different initial opinion level in the general group: $(a,\ d,\ g)$ $x_1(0) = 0.65$, $(b,\ e,\ h)$ $x_1(0) = 0.67$, and $(c,\ f,\ i)$ $x_1(0) = 0.69$. The small panels on the left show the time-varying homophily levels for each α.

back to a high level (Figure 4.4).

Monotonic changes in homophily levels have a smaller impact on the general group than on the minority groups. Due to the relative size of each group, within-group influence always dominates the dynamics in the general group. When homophily increases over time, this dominance is reinforced, further amplifying the role of internal imitation. When homophily decreases, cross-group influence becomes more significant. However, since minority groups tend to converge to the dynamics of the general group under low homophily, the overall trajectory still reflects that of the general group. That said, the general group can still undergo noticeable changes in the early stages of monotonic homophily. For example, when homophily strengthens faster, vaccine support in the general group declines more slowly, or may not decline at all (Figure 4.3b, e, and h).

4.4.3 Time evolution of homophily-driven opinion polarization

The existence of multi-strategy equilibria revealed in the previous sections highlights that opinion polarization can persist in the long term, regardless of the presence of disease. Here, we measure opinion polarization by looking at how much the minority groups deviate from the general group. We define the Minority-Majority Divergence (MMD) at time t as the population-weighted absolute deviation of minority groups' pro-vaccine percentages from the majority group's pro-vaccine percentage:

$$\mathrm{MMD}(t) = \frac{N_2|x_2 - x_1| + N_3|x_3 - x_1|}{N_2 + N_3},$$

and plot the time evolution under different basic reproduction numbers R_0, shown in Figure 4.5.

When $R_0 < 1$, the polarization in vaccine sentiments begins to diminish under two conditions: (i) when both h_2 and h_3 are small, or (ii) when h_2 is large (Figure 4.5*a*–*e*). However, these two scenarios lead to different consensus outcomes. In the first case, when homophily is weak in both minority groups, all three groups converge to a pro-vaccine consensus. In the second case, when the pro-vaccine group has high homophily, the system reaches an anti-vaccine consensus (Figure 4.6*a*). It is important to note that when $R_0 < 1$, each infectious individual transmits the disease to fewer than one other person on average, causing the disease to eventually die out. As a result, the behavioral dynamics resemble those of a disease-free system, as illustrated in Figure 4.2*a*, but it is overall more complex.

As the basic reproduction number increases beyond the bifurcation point ($R_0 = 1$) to $R_0 = 4$, we do not observe major qualitative differences in the time evolution or the final state of MMD and overall vaccine support over the parameter plane (Figure 4.5*f*–*j*, Figure 4.6*b*). Due to the occurrence of large-scale outbreaks, the pro-vaccine percentage in each group, after an initial decline that occurs particularly when h_2 is large, experiences varying degrees of increase, leading to periodic and moderate increases in polarization (Figure 4.5*g* and *h*).

When R_0 increases further to around 8, a new region of depolarization emerges on the parameter plane, specifically where h_2 ranges from 0.6 to 0.8, and h_3 remains below 0.6 (Figure 4.5*k*–*o*). In these scenarios, the higher transmissibility of disease leads to an earlier initial outbreak, which raises the pro-vaccine percentage across all three groups, eventually driving the system toward near-universal vaccine support. As a result, the pro-vaccine consensus region expands, while the anti-vaccine consensus region shrinks in Figure 4.6*c*.

Finally, for a very large basic reproduction number $R_0 = 10$, such as that associated with measles, the persistence and disappearance of polarization are almost entirely

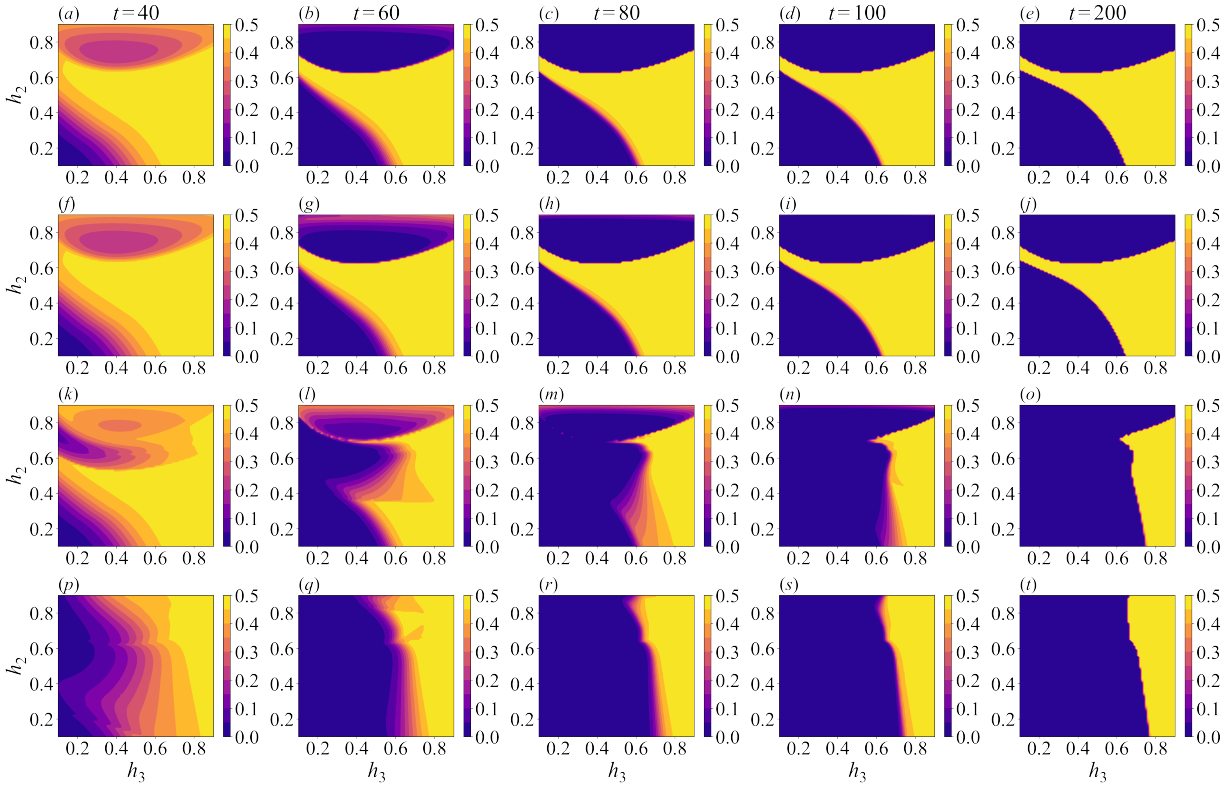


Figure 4.5: Time evolution of Minority-Majority Divergence under varying basic reproduction numbers over the h_2–h_3 parameter plane. Each row corresponds to a different basic reproduction number: $(a$–$e)$ $R_0 = 0.8$, $(f$–$j)$ $R_0 = 4$, $(k$–$o)$ $R_0 = 8$, and $(p$–$t)$ $R_0 = 10$. Columns represent different time points: $t = 40$, 60, 80, 100, and 200, respectively. The initial pro-vaccine percentage in the general group $x_1(0) = 0.67$, and the initial Minority-Majority Divergence $\mathrm{MMD}(0) = 0.375$.

governed by h_3 (Figure 4.5p–t), and the anti-vaccine consensus will no longer exist (Figure 4.6d).

4.5 Discussion

Social heterogeneity should not be overlooked in infectious disease modeling. It plays a crucial role not only on social structure but also on the biological level, where it indirectly influences disease dynamics by shaping individuals' sentiments on vaccination. Homophily

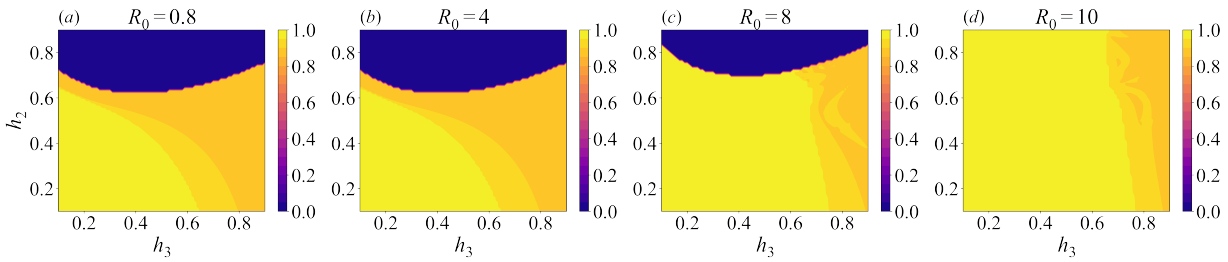


Figure 4.6: Pro-vaccine percentage in the entire population at $t = 200$ under varying basic reproduction numbers over the h_2–h_3 parameter plane, corresponding to Figure 4.5*e*, *j*, *o*, and *t*.

further amplifies these effects by forming clusters among individuals with similar opinions. In this study, we introduce a multi-group vaccination game model to explore how population heterogeneity and homophily jointly affect disease and behavioral dynamics. We focus on a simple setup with a majority group that represents the general public, and two minority groups of strongly pro-vaccine and anti-vaccine people.

Our results reveal that mixed-strategy equilibria commonly emerge at the group level, and population-level opinion polarization can persist regardless of whether a disease is present, suggesting that social structure alone is enough to maintain divided vaccination behavior. We also show that time-varying homophily strongly affects the extent of polarization. These findings highlight the importance of evolving social dynamics in shaping long-term public attitudes toward vaccination.

The emergence of a group-level mixed-strategy equilibrium, where both pro-vaccine and anti-vaccine individuals persist within the same group, can be understood by examining the within-group and cross-group interactions. With partial homophily, imitation dynamics are shaped by both within-group and cross-group payoff comparisons, making it possible to reach a balance where these opposing effects cancel each other out. For example, within-group interactions favor pro-vaccine strategies in a group where most individuals are already pro-vaccine, while cross-group interactions may still favor anti-vaccine strategies. In such a case, anti-vaccine individuals imitate pro-vaccine peers within the group, while those pro-vaccine members may be influenced by anti-vaccine individuals from other groups and thus change their strategy. Meanwhile, an anti-vaccine-dominant group may exhibit the opposite pattern. At the population level, such group-specific dynamics give rise to persistent opinion polarization, which means that a highly pro-vaccine group can coexist with a highly anti-vaccine group in the long run because high homophily isolates minority groups from being fully assimilated into the majority's stance.

In a multi-group vaccination game, the dual role of homophily is more prominently exhibited. Although high homophily reduces communication between groups, it can be beneficial when cross-group interaction proves harmful to the interests of a group. For example, a highly pro-vaccine group may benefit from stronger within-group ties, as frequent interactions with a group having a lower pro-vaccine percentage can slow down the growth in vaccine support. Low homophily facilitates broader information exchange and increases the chances of pro-vaccine strategies spreading across groups, but it also allows anti-vaccine strategies to propagate more easily, which amplifies the negative effect of highly anti-vaccine groups. From a population-level perspective, having polarization across groups may be more desirable than convergence to an anti-vaccine consensus.

The general group with moderate homophily and vaccine support near the critical threshold initially behaves like a swing group, influenced by the pro- and anti-vaccine minorities. As the group converges to a single strategy, it gradually becomes the dominant factor shaping the dynamics of both minority groups. An asymmetry emerges in how the pro- and anti-vaccine groups respond to this influence when homophily in minority groups is high. Specifically, the highly anti-vaccine group can sustain a low vaccine support even when the broader population moves toward a pro-vaccine consensus. In contrast, when the general group's pro-vaccine percentage declines toward zero, the highly pro-vaccine group often fails to maintain its stance, especially when the disease is not severe. In such cases, the impact of the disease can be neglected, and the perceived risk of the vaccine makes vaccination a less attractive strategy. Therefore, if cross-group imitation favors anti-vaccine behavior, a strong social norm is required for the pro-vaccine group to ensure that within-group imitation outweighs the external pressure of vaccine refusal.

At first glance, it may seem surprising that opinion polarization can persist at certain levels of homophily, even as the disease becomes more severe. One might expect that a higher transmission rate would elevate the perceived risk of infection and encourage vaccination across all groups. However, this persistence can be explained by three reasons. First, in the anti-vaccine group with high homophily, behavioral dynamics are largely driven by within-group imitation, which reinforces existing anti-vaccine strategies. Second, although the infection risk increases with higher transmission rates, it may still be insufficient to override the combined effects of the vaccine risk and social norms. Third, higher transmission rates prompt faster vaccine uptake in the general and pro-vaccine groups, leading to widespread immunity in the entire population, which lowers the number of infectious individuals and further reduces the urgency of vaccination within the anti-vaccine group.

It should be noted that the parametric pair (h_2, h_3) induces a first-order "phase transition"-like behavior (for the legitimacy of the use of the term "phase transition" in finite-dimensional models, see [208]) in the steady state of $x(t)$ in Figure 4.2. When $x_1(0) = 0.67$,

the regions corresponding to the low and high "phases" are separated by a non-monotone curve in the h_2–h_3 parameter plane; when $x_1(0) = 0.69$, the transition still exists, but it is much weaker and the separation curve is monotone. Moreover, considering MMD(t) as an order parameter, we observe first-order phase-transition-like behavior at late times, such as $t = 200$ (Figure 4.5). When R_0 is small, the parameter plane is divided into three regions with non-monotonic borders (Figure 4.5*e* and *j*), instead of two as in Figure 4.2*a*. Two of these regions correspond to "high phases" and the third corresponds to a "low phase". For earlier times like $t = 80$, there is a coexistence of first- and second-order transitions. Interestingly, for larger values of R_0, there is a qualitative change in the order parameter density plot: it is divided into only two regions, one corresponding to the high phase and the other to the low phase (which is now smaller in size). It is also worth noting that the border between the two regions is far less smooth (Figure 4.5*o* and *t*). Critical transitions in dynamical systems and phase transitions in statistical physics share some similar features, but a rigorous theoretical connection between them is still lacking. Future research could focus on clarifying this link, which may lead to a deeper understanding of early warning signals in coupled behavior-disease systems [106].

One notable limitation of our model is that it does not account for heterogeneity at the biological level, so the feedback loop between disease dynamics and vaccination behavior may be misrepresented. Long-standing research on age-specific contact patterns relevant to the transmission of childhood infectious diseases has shown that such interactions are far from homogeneous (i.e., random), rather they are typically characterized by strong degrees of heterogeneity and homophily (i.e., assortativeness) [7, 115, 166, 81]. Recent work estimated that the basic reproduction number for measles can increase by 70% if heterogeneity is taken into account [101]. Therefore, realistic coupled behavior-disease models should account for a double heterogeneity in both social contacts and opinion dynamics. However, we must notice that when modeling childhood diseases in a coupled behavior-disease model, the agents driving vaccination decisions are distinct from those involved in disease transmission. Therefore, when incorporating heterogeneity on both biological and social layers, these layers should be properly aligned, which means each group of individuals on the biological layer should be associated with their parents on the social layer, who make the vaccination decisions on their behalf. Additionally, our model does not include stochastic effects or spatial mobility of agents, both of which may further influence disease spread and behavioral responses in real-world settings.

In this study, we develop a multi-group model for an arbitrary number of groups, but only focus on simulating three-group dynamics to strike a balance between computational complexity and the clarity of resulting insights. As additional pro- and anti-vaccine groups with different levels of vaccine support are considered, they not only collectively influence

the general public but also affect one another, leading to more complex dynamics. Moreover, dividing the population into majority and extreme minority groups is primarily a conceptual approach. In practice, the choice of grouping should depend on the available data. For example, individuals can be grouped by geographic regions, as vaccination rates often show spatial clustering [78, 145, 39]. Geographic grouping also captures homophily, since people living in the same area are more likely to interact with each other. In such cases, a large general group may not exist, and group sizes may be more evenly distributed.

In our model, we only considered long-term exogenous (i.e., non-epidemiological) changes in homophily to identify stable influences of homophily trends on vaccination behavior. However, previous research has shown that exogenous conditions can cause homophily to evolve quite rapidly [144], which could affect infection trajectories in the short term. In relation to this, an example might be seasonal influenza epidemics over successive seasons. Additionally, homophily might also evolve endogenously. For instance, both epidemic outbreaks and sudden rumors about vaccine adverse events could trigger faster changes in homophily (e.g., by raising stronger concerns about vaccine effectiveness or safety even in the majority group).

Our findings highlight that homophily, whether static or dynamic, plays a crucial role in shaping long-term vaccination behavior and opinion polarization. While lowering homophily can reduce polarization and help the spread of pro-vaccine strategies, it also unintentionally creates conditions for the dissemination of anti-vaccine strategies. As a result, from a public health perspective, simply removing communication barriers is not enough; effective control over the content being spread is equally important. In summary, our multi-group framework offers greater flexibility for simulating the dynamics of vaccination behavior across diverse population structures. Future work could validate the model with empirical data and extend the framework by incorporating heterogeneous transmission dynamics.

Chapter 5

Conclusion

This thesis aimed to advance our understanding of the dynamics of vaccination behavior in structured populations, especially in the context of childhood diseases such as measles. It brought together three research articles that collectively addressed two main knowledge gaps: (i) How do population heterogeneity and homophily influence vaccination behavior and disease dynamics? (ii) Can signals of declining vaccine sentiment be identified from high-frequency, noisy data from sources such as social media by deep learning models?

To investigate these problems, we integrated mathematical modeling and machine learning techniques in successive stages. We first extended existing coupled behavior-disease models by incorporating imitation dynamics between two social groups (social media users and non-users), highlighting how homophily and online randomness shape vaccination outcomes. Motivated by the work on combining mechanistic and data-driven approaches [14] and building on our modeling framework, we then designed two deep learning classifiers trained on model-generated social media data to detect early warning signals and predict outbreak risk using empirical data. Finally, we generalized our model to a flexible multi-group setting, with a focus on the emergence and persistence of opinion polarization.

5.1 Summary of findings

Homophily refers to the tendency for individuals to associate with others who are similar to themselves. In this thesis, it is reflected in its role in facilitating within-group interaction while hindering cross-group communication. At its core, homophily functions as a neutral structural force in social networks, which means that it does not inherently promote or

suppress any particular type of information, whether accurate or misleading. However, its impact can be dual in nature. Depending on the content circulating within a group, homophily may either reinforce positive outcomes, such as the spread of scientifically valid information, or amplify harmful effects, such as the persistence of false beliefs. Therefore, while homophily itself does not carry value judgment, the consequences it produces can be significantly affected by the informational environment in which it operates.

In infectious disease modeling, analysis is often performed around equilibrium points because these points provide critical insights into the long-term behavior and stability of disease dynamics. Typically, epidemic models feature two primary types of equilibria: a disease-free equilibrium, where the infection is eradicated from the population, and an endemic equilibrium, where the disease persists at a stable prevalence. When vaccination behavior is coupled with disease dynamics, the number of equilibrium points increases. The differential equation governing the time evolution of the vaccine support can yield multiple steady states, such as fully pro-vaccine, fully anti-vaccine, and partially pro-vaccine (i.e., partially anti-vaccine). These behavioral equilibria can combine with disease-related equilibria to create more diverse long-term outcomes. The system becomes more complex in multi-group models, where the population structure introduces within- and cross-group imitation processes. In such settings, the system can settle into additional equilibria at which different groups exhibit distinct strategy profiles. These new equilibria often emerge from a balance of opposing imitation directions.

Thus, the effects of homophily on coupled behavior-disease dynamics crucially depend on the equilibrium to which the system converges. In Chapter 2, we examined the disease-free equilibrium where both groups adopt a fully pro-vaccine stance. On the one hand, this represents the desired long-term outcome and closely reflects the state of many real-world systems, characterized by low infection prevalence and high vaccination coverage. On the other hand, in multi-group models, this equilibrium is also among the easiest to identify analytically, compared to others involving partial vaccine uptake or asymmetric group behavior. In this scenario, homophily reinforces the pro-vaccine norm in each group, but simultaneously hinders the cross-group spread of pro-vaccine strategies. In Chapter 4, our focus shifted to other possible equilibrium points, including those with partial or complete vaccine refusal. Specifically, we initialized each group near distinct equilibrium points derived under the assumption of homogeneous mixing, ranging from fully pro-vaccine to partially pro-vaccine and fully anti-vaccine stances. Since cross-group interactions are permitted, combinations of such equilibria do not necessarily correspond to an equilibrium of the coupled multi-group system. For highly pro-vaccine groups, homophily promotes the spread of pro-vaccine strategies within the group, while impeding the inflow of anti-vaccine influence from other groups. For highly anti-vaccine groups, the effect is reversed. In

partially pro-vaccine groups, the effect of homophily depends on the relative level of vaccine support inside and outside the group. When the internal support is higher, homophily has a beneficial effect. The opposite holds when the internal support is lower.

In Chapters 2 and 4, homophily is implemented in the model in different ways. In Chapter 2, the parameter h acts like a valve controlling the flow of out-group information. Under complete homophily, the valve is closed, and only within-group imitation occurs. As homophily decreases, the valve opens, allowing more information from other groups to enter. This approach works well when the number of groups is small and when group sizes and vaccine support levels are relatively balanced (e.g., nearly fully pro-vaccine in two groups). As the number of groups increases, adjusting h alone to control cross-group interaction can distort the dynamics, making them driven more by the amount of imitation than by homophily itself. Moreover, if group sizes and vaccine support levels differ greatly, the influence of a very small group could be amplified, potentially offsetting that of a much larger group, which does not align with real-world dynamics. In Chapter 4, we improved the approach by assuming that the rate at which individuals encounter others' opinions is proportional to the relative group size, which means that strategies from larger groups are more likely to be seen–a matter of probability. Homophily now controls how individuals allocate attention to these opinions. Higher homophily levels lead to a stronger preference for in-group information, while lower levels reduce this bias, with the extreme case being no distinction between in- and out-group information. This mechanism resembles a parallel circuit, where all information channels are open, but attention (like current) is distributed across them based on individual preferences (like resistances).

On the behavioral layer of coupled behavior-disease models, heterogeneity can arise from both direct and indirect sources. Direct heterogeneity refers to observable differences in vaccine sentiments across groups. Indirect heterogeneity stems from social or structural factors that shape how individuals access and process information. For example, social media users and non-users can have different levels of exposure to vaccine-related content, even if their support for vaccines could be similar. Chapter 2 captures indirect heterogeneity by applying a larger noise amplitude for social media users than non-users in the stochastic model. Social media users are typically exposed to more rapid and diverse information flows, which can lead to faster attitude shifts and more dramatic changes in overall vaccine support. In contrast, non-users show more gradual behavioral changes. From a dynamical systems perspective, stronger fluctuations may push the system across the boundaries of a basin of attraction, resulting in transitions to alternative stable states. This transition is also evident in our results: under widespread vaccine scare, vaccine collapse tends to occur earlier among social media users. Chapter 4 focuses on direct heterogeneity by assigning different initial levels of vaccine support to each group. This setup creates an initial state

of opinion polarization: while both pro- and anti-vaccine individuals exist in the population, they are not evenly distributed but clustered into groups. The goal is not to analyze consensus equilibrium points (i.e., fully pro- or anti-vaccine states in all groups), but rather to explore the implications of preexisting differences in vaccination behavior, such as those rooted in cultural norms or historical trust in vaccination. The analysis shows that the initial opinion polarization can persist when homophily is high, but eventually gives way to consensus when homophily is low. Notably, this approach does not include stochasticity, so behavioral dynamics are entirely driven by the deterministic model.

Real-time monitoring of vaccination behavior is critical for forecasting outbreak risk, as declines in vaccine uptake often precede increases in disease incidence. However, it faces two main challenges. First, vaccination rate data are often temporally coarse, so higher-frequency data sources are required for timely feedback. Second, once reliable data are available, appropriate methods are needed to detect underlying behavioral trends. The project spanning Chapters 2 and 3 is inspired by two previous studies. The first identifies EWS, specifically increased variance and lag-1 autocorrelation, in vaccine-related Tweets before a major measles outbreak in the US [175]. The second develops a deep learning classifier based on the CNN-LSTM framework to detect whether a system is approaching a tipping point and to predict the type of bifurcation from time series data [38].

Building on these studies, we developed a hybrid approach that integrates mechanistic modeling with data-driven techniques to detect EWS of vaccine sentiment collapse from social media data. Chapter 2 introduces a stochastic coupled behavior-disease model that captures imitation dynamics among social media users under the influence of online information, incorporating fat-tailed Lévy noise to reflect the extreme fluctuations often observed in social media data. This model not only simulates how shifts in public vaccine sentiment may precede outbreaks but also generates synthetic time series that mimic real-world vaccine-related discussions. In Chapter 3, we used these model-generated data to train two deep learning classifiers. Besides the well-established CNN-LSTM framework, we also built a simple model based on the ResNet architecture. These DL classifiers outperform traditional statistical CSD indicators like variance and lag-1 autocorrelation, particularly when applied to noisy time series with non-Gaussian noise and to transformed data that are further perturbed from the outputs of differential equation systems. We also showed that the DL classifiers retain predictive performance on empirical social media data, despite being trained purely on simulated data. The results highlight the promise of hybrid approaches when real-world training data are limited.

5.2 Future work

Mechanistic models aim to explain real-world phenomena by representing the underlying processes of a system using mathematical structures. They serve as tools for understanding not just what is happening, but why it is happening. In behavioral epidemiology, coupled behavior-disease models explain the feedback loop between disease transmission and vaccination behavior [93, 152, 229, 223]. Since the concept of the "vaccination game" was introduced 20 years ago, researchers have progressively incorporated factors such as vaccine scare and social norms to capture the complexity of human behavioral dynamics [20, 21, 173]. This thesis contributes to this line of work by incorporating population structure on the behavioral layer into the coupled model, with an emphasis on the impact of homophily. Future work could build on our work in several directions. First, models should account for both direct heterogeneity (i.e., differences in vaccine support) and indirect heterogeneity (e.g., differences in information exposure), since they often coexist and affect each other [75]. Heterogeneity on the behavioral layer should be considered together with heterogeneity on the disease layer, such as contact patterns, to reflect the multi-layered nature of real-world epidemics [228]. Second, more attention should be paid to validating mechanistic models against empirical data, since their usefulness for public health depends on how well they capture observed behavioral and epidemiological patterns [92, 233].

While mechanistic models have mainly been used to control outbreaks once they have started, future work should consider their role in preventing outbreaks before they occur. Transitioning from reactive to proactive modeling requires estimating outbreak risk by accounting for complex behavioral and environmental factors, a task that mechanistic models often struggle with because they rely on predefined assumptions and precise parameterization [180, 119]. In contrast, data-driven models offer greater flexibility and adaptability by learning directly from diverse and noisy real-world data without relying on strict assumptions. However, developing reliable data-driven models requires large volumes of high-quality data, which are often difficult or even impossible to obtain in practice [4]. Hybrid modeling approaches that integrate mechanistic insights with data-driven techniques hold significant promise [14, 244]. Future work should consider coupled behavior-disease models not only as tools for understanding disease and behavioral dynamics, but also as generators of synthetic data to support the training of data-driven models [65, 99]. Improved coupled behavior-disease models should not be used solely to produce a single accurate trajectory. Instead, they should be employed to create a wide range of plausible scenarios by varying parameters and introducing stochasticity, thus improving the robustness of data-driven models [5].

Enhancing the predictive ability of such hybrid approaches also requires identifying

and improving data sources that reflect real-time changes in vaccine-related behaviors and sentiments. Although social media offer high-frequency data on public sentiment, they come with limitations such as demographic bias and high noise levels. Future work should explore a broader range of data sources that may reveal latent vaccine sentiment, such as search engine trends, online forums, and mobility data [16]. At the same time, advances are needed in sentiment analysis, natural language processing, and geocoding to extract more accurate information. Mechanistic models should also explain the structural relationship between these data sources and behavioral dynamics, enabling them to generate training data for data-driven models [224].

5.3 Concluding remarks

This thesis contributes to the growing body of work at the intersection of behavioral epidemiology and data-driven dynamical systems by investigating how population heterogeneity and homophily shape vaccination behavior, and by exploring early warning signals of vaccine sentiment shifts using data-driven approaches. Through a sequence of three studies, we developed and analyzed coupled behavior-disease models that incorporate imitation dynamics within structured populations, which then serve as generators of training data for deep learning classifiers to detect critical transitions in vaccine sentiment. Our results reveal the role of homophily in shaping long-term opinion polarization, and demonstrate the feasibility of using social media data to reflect real-time changes in vaccine sentiment. By integrating mechanistic models and data-driven techniques, we showed that hybrid modeling approaches can estimate the risk of outbreaks from online behavioral signals. Future research should focus on expanding couple behavior-disease models to include multi-layer heterogeneity and using them to support the training of data-driven models, paving the way for hybrid approaches to be more widely applied in public health interventions.

References

[1] G. O. Agaba, Y. N. Kyrychko, and K. B. Blyuss. Mathematical model for the impact of awareness on the dynamics of infectious diseases. *Math. Biosci.*, 286:22–30, 2017.

[2] A. E. Aiello, A. Renson, and P. N. Zivich. Social media- and Internet-based disease surveillance for public health. *Annu. Rev. Public Health*, 41:101–118, 2020.

[3] S. Ajovalasit, V. M. Dorgali, A. Mazza, A. d'Onofrio, and P. Manfredi. Evidence of disorientation towards immunization on online social media after contrasting political communication on vaccines. Results from an analysis of Twitter data in Italy. *PLoS One*, 16:e0253569, 2021.

[4] T. Alamo, D. G. Reina, P. Millán Gata, V. M. Preciado, and G. Giordano. Data-driven methods for present and future pandemics: Monitoring, modelling and managing. *Annu. Rev. Control*, 52:448–464, 2021.

[5] M. Alber, A. Buganza Tepole, W. R. Cannon, S. De, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W. W. Lytton, P. Perdikaris, L. Petzold, and E. Kuhl. Integrating machine learning and multiscale modeling-perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Digit. Med.*, 2:115, 2019.

[6] W. Alotaibi, F. Alomary, and R. Mokni. COVID-19 vaccine rejection causes based on Twitter people's opinions analysis using deep learning. *Soc. Netw. Anal. Min.*, 13:62, 2023.

[7] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1991.

[8] E. Aramaki, S. Maskawa, and M. Morita. Twitter catches the flu: detecting influenza epidemics using Twitter. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1568–1576. Association for Computational Linguistics, 2011.

[9] K. M. Ariful Kabir, K. Kuga, and J. Tanimoto. The impact of information spreading on epidemic vaccination game dynamics in a heterogeneous complex network- A theoretical approach. *Chaos Solitons Fractals*, 132:109548, 2020.

[10] D. I. Armstrong McKay, A. Staal, J. F. Abrams, R. Winkelmann, B. Sakschewski, S. Loriani, I. Fetzer, S. E. Cornell, J. Rockstrőm, and T. M. Lenton. Exceeding 1.5°C global warming could trigger multiple climate tipping points. *Science*, 377:eabn7950, 2022.

[11] S. Ashkenazi, G. Livni, A. Klein, N. Kremer, A. Havlin, and O. Berkowitz. The relationship between parental source of information and knowledge about measles / measles vaccine and vaccine hesitancy. *Vaccine*, 38:7292–7298, 2020.

[12] İ. Aygün, B. Kaya, and M. Kaya. Aspect based Twitter sentiment analysis on vaccination and vaccine types in COVID-19 Pandemic with deep learning. *IEEE J. Biomed. Health Inform.*, 26:2360–2369, 2022.

[13] L. R. Baghdadi, M. M. Hassounah, A. Younis, H. I. Al Suwaidan, and R. Al Khalifah. Caregivers' Sources of Information About Immunization as Predictors of Delayed Childhood Vaccinations in Saudi Arabia During the COVID-19 Pandemic: A Cross-Sectional Questionnaire Study. *Risk Manag. Healthc. Policy*, 14:3541–3550, 2021.

[14] R. E. Baker, J.-M. Peña, J. Jayamohan, and A. Jérusalem. Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biol. Lett.*, 14:20170660, 2018.

[15] E. Bakshy, S. Messing, and L. A. Adamic. Political science. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348:1130–1132, 2015.

[16] S. Bansal, G. Chowell, L. Simonsen, A. Vespignani, and C. Viboud. Big Data for Infectious Disease Surveillance and Modeling. *J. Infect. Dis.*, 214:S375–S379, 2016.

[17] S. Bansal, B. T. Grenfell, and L. A. Meyers. When individual behaviour matters: homogeneous and network models in epidemiology. *J. R. Soc. Interface*, 4:879–891, 2007.

[18] A. Barbagallo and M. G. Cojocaru. Dynamic vaccination games and variational inequalities on time-dependent sets. *J. Biol. Dyn.*, 4:539–558, 2010.

[19] C. Bauch, A. d'Onofrio, and P. Manfredi. *Behavioral Epidemiology of Infectious Diseases: An Overview*, chapter 1, pages 1–19. Springer New York, 2012.

[20] C. T. Bauch. Imitation dynamics predict vaccinating behaviour. *Proc. R. Soc. B Biol. Sci.*, 272(1573):1669–1675, 2005.

[21] C. T. Bauch and S. Bhattacharyya. Evolutionary Game Theory and Social Learning Can Determine How Vaccine Scares Unfold. *PLoS Comput. Biol.*, 8(4):e1002452, 2012.

[22] C. T. Bauch and D. J. D. Earn. Vaccination and the theory of games. *Proc. Natl. Acad. Sci. U.S.A.*, 101(36):13391–13394, 2004.

[23] C. T. Bauch and A. P. Galvani. Social Factors in Epidemiology. *Science*, 342:47–49, 2013.

[24] C. T. Bauch, A. P. Galvani, and D. J. D. Earn. Group interest versus self-interest in smallpox vaccination policy. *Proc. Natl. Acad. Sci. U. S. A.*, 100:10564–10567, 2003.

[25] O. Benecke and S. E. DeYoung. Anti-Vaccine Decision-Making and Measles Resurgence in the United States. *Glob. Pediatr. Health*, 6:2333794X1986294, 2019.

[26] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.*, 5:157–166, 1994.

[27] K. Betz, F. Fu, and N. Masuda. Evolutionary Game Dynamics with Environmental Feedback in a Network with Two Communities. *Bull. Math. Biol.*, 86:84, 2024.

[28] S. Bhattacharyya and C. T. Bauch. "Wait and see" vaccinating behaviour during a pandemic: A game theoretic analysis. *Vaccine*, 29(33):5519–5525, 2011.

[29] S. Bidari and W. Yang. Global resurgence of measles in the vaccination era and influencing factors. *Int. J. Infect. Dis.*, 147:107189, 2024.

[30] N. Boers. Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record. *Nat. Commun.*, 9:2556, 2018.

[31] F. Brauer, P. van den Driessche, and J. Wu. *Mathematical Epidemiology*. Springer, 2008.

[32] T. S. Brett, J. M. Drake, and P. Rohani. Anticipating the emergence of infectious diseases. *J. R. Soc. Interface*, 14(132):20170115, 2017.

[33] T. S. Brett, E. B. O'Dea, É. Marty, P. B. Miller, A. W. Park, J. M. Drake, and P. Rohani. Anticipating epidemic transitions with imperfect data. *PLoS Comput. Biol.*, 14(6):e1006204, 2018.

[34] B. Buonomo, N. Chitnis, and A. d'Onofrio. Seasonality in epidemic models: a literature review. *Ric. Mat.*, 67:7–25, 2018.

[35] D. C. Burgess, M. A. Burgess, and J. Leask. The MMR vaccination and autism controversy in United Kingdom 1998-2005: inevitable community outrage or a failure of risk communication? *Vaccine*, 24:3921–3928, 2006.

[36] T. M. Bury. ewstools: A Python package for early warning signals of bifurcations in time series data. *J. Open Source Softw.*, 8(82):5038, 2023.

[37] T. M. Bury, D. Dylewsky, C. T. Bauch, M. Anand, L. Glass, A. Shrier, and G. Bub. Predicting discrete-time bifurcations with deep learning. *Nat. Commun.*, 14:6331, 2023.

[38] T. M. Bury, R. I. Sujith, I. Pavithran, M. Scheffer, T. M. Lenton, M. Anand, and C. T. Bauch. Deep learning for early warning signals of tipping points. *Proc. Natl. Acad. Sci. U.S.A.*, 118(39):e2106140118, 2021.

[39] J. Cadena, D. Falcone, A. Marathe, and A. Vullikanti. Discovery of under immunized spatial clusters using network scan statistics. *BMC Med. Inform. Decis. Mak.*, 19:28, 2019.

[40] California Department of Public Health. School Immunizations in Kindergarten by Academic Year, 2020.

[41] A. Carlson, M. Riethman, P. Gastañaduy, A. Lee, J. Leung, M. Holshue, C. DeBolt, and A. Melnick. Notes from the Field: Community Outbreak of Measles — Clark County, Washington, 2018–2019. *MMWR. Morb. Mortal. Wkly. Rep.*, 68(19):446–447, 2019.

[42] Centers for Disease Control and Prevention. About the National Immunization Surveys (NIS), 2018.

[43] Centers for Disease Control and Prevention. Measles History, 2020.

[44] Centers for Disease Control and Prevention. Measles, mumps, and rubella (MMR) vaccination: what everyone should know, 2021.

[45] Centers for Disease Control and Prevention. Explaining how vaccines work, 2025.

[46] D. Centola. The Spread of Behavior in an Online Social Network Experiment. *Science*, 329(5996):1194–1197, 2010.

[47] D. Centola. An Experimental Study of Homophily in the Adoption of Health Behavior. *Science*, 334:1269–1272, 2011.

[48] S. L. Chang, M. Piraveenan, P. Pattison, and M. Prokopenko. Game theoretic modelling of infectious disease dynamics and intervention methods: a review. *J. Biol. Dyn.*, 14:57–89, 2020.

[49] G. B. Chapman, M. Li, J. Vietri, Y. Ibuka, D. Thomas, H. Yoon, and A. P. Galvani. Using Game Theory to Examine Incentives in Influenza Vaccination Behavior. *Psychol. Sci.*, 23:1008–1015, 2012.

[50] X. Chen and F. Fu. Imperfect vaccine and hysteresis. *Proc. R. Soc. B Biol. Sci.*, 286(1894):20182406, 2019.

[51] R. B. Cialdini. *Influence: The Psychology of Persuasion*. Harper Business, 1984.

[52] R. B. Cialdini, R. R. Reno, and C. A. Kallgren. A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Personality and Social Psychology*, 58(6):1015–1026, 1990.

[53] M. Cinelli, G. De Francisci Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini. The echo chamber effect on social media. *Proc. Natl. Acad. Sci. U.S.A.*, 118:e2023301118, 2021.

[54] M. G. Cojocaru. Dynamic equilibria of group vaccination strategies in a heterogeneous population. *J. Glob. Optim.*, 40:51–63, 2008.

[55] M. G. Cojocaru, C. T. Bauch, and M. D. Johnston. Dynamics of Vaccination Strategies via Projected Dynamical Systems. *Bull. Math. Biol.*, 69:1453–1476, 2007.

[56] N. D. Cook. *Born to Die: Disease and New World Conquest, 1492–1650*. Cambridge University Press, 1998.

[57] A. W. Crosby. *The Columbian Exchange: Biological and Cultural Consequences of 1492*. Greenwood Publishing Company, 1972.

[58] A. W. Crosby. *Ecological Imperialism: The Biological Expansion of Europe, 900–1900*. Cambridge University Press, 1986.

[59] J. Cui, Y. Sun, and H. Zhu. The impact of media on the control of infectious diseases. *J. Dyn. Diff. Equat.*, 20:31–53, 2007.

[60] F. T. Cutts, P. Claquin, M. C. Danovaro-Holliday, and D. A. Rhoda. Monitoring vaccination coverage: defining the role of surveys. *Vaccine*, 34:4103–4109, 2016.

[61] A. Dabbagh, R. L. Laws, C. Steulet, L. Dumolard, M. N. Mulders, K. Kretsinger, J. P. Alexander, P. A. Rota, and J. L. Goodson. Progress Toward Regional Measles Elimination — Worldwide, 2000–2017. *Morb. Mortal. Wkly. Rep.*, 67(47):1323–1329, 2018.

[62] V. Dakos, C. A. Boulton, J. E. Buxton, J. F. Abrams, B. Arellano-Nava, D. I. Armstrong McKay, S. Bathiany, L. Blaschke, N. Boers, D. Dylewsky, C. López-Martínez, I. Parry, P. Ritchie, B. van der Bolt, L. van der Laan, E. Weinans, and S. Kéfi. Tipping point detection and early warnings in climate, ecological, and human systems. *Earth Syst. Dynam.*, 15:1117–1135, 2024.

[63] V. Dakos, S. R. Carpenter, W. A. Brock, A. M. Ellison, V. Guttal, A. R. Ives, S. Kéfi, V. Livina, D. A. Seekell, E. H. van Nes, and M. Scheffer. Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data. *PLoS One*, 7(7):e41010, 2012.

[64] V. Dakos, M. Scheffer, E. H. van Nes, V. Brovkin, V. Petoukhov, and H. Held. Slowing down as an early warning signal for abrupt climate change. *Proc. Natl. Acad. Sci. U.S.A.*, 105:14308–14312, 2008.

[65] C. M. de Melo, A. Torralba, L. Guibas, J. DiCarlo, R. Chellappa, and J. Hodgins. Next-generation deep learning based on simulators and synthetic data. *Trends Cogn. Sci.*, 26:174–187, 2022.

[66] S. Deb, S. Sidheekh, C. F. Clements, N. C. Krishnan, and P. S. Dutta. Machine learning methods trained on simple models can predict critical transitions in complex natural systems. *R. Soc. Open Sci.*, 9:211475, 2022.

[67] A. N. Desai and M. S. Majumder. What Is Herd Immunity? *JAMA*, 324(20):2113, 2020.

[68] J. M. Diamond. *Guns, Germs, and Steel: The Fates of Human Societies*. W. W. Norton & Co., 1997.

[69] R. C. Dicker, F. Coronado, D. Koo, and R. G. Parrish. *Principles of epidemiology in public health practice; an introduction to applied epidemiology and biostatistics*. Centers for Disease Control and Prevention (CDC), 2006.

[70] S. Dinesh and M. Odabaş. 8 facts about Americans and Twitter as it rebrands to X, 2023.

[71] M. K. Doll and J. W. Correira. Revisiting the 2014-15 Disneyland measles outbreak and its influence on pediatric vaccinations. *Hum. Vaccin. Immunother.*, 17:4210–4215, 2021.

[72] A. d'Onofrio, P. Manfredi, and P. Poletti. The Interplay of Public Intervention and Private Choices in Determining the Outcome of Vaccination Programmes. *PLoS One*, 7(10):e45653, 2012.

[73] A. d'Onofrio, P. Manfredi, and E. Salinelli. Vaccinating behaviour, information, and the dynamics of SIR vaccine preventable diseases. *Theor. Popul. Biol.*, 71:301–317, 2007.

[74] A. A. Dror, N. Eisenbach, S Taiber, N. G. Morozov, M. Mizrachi, A. Zigron, S. Srouji, and E. Sela. Vaccine hesitancy: the next challenge in the fight against COVID-19. *Eur. J. Epidemiol.*, 35:775–779, 2020.

[75] A. G. Dunn, D. Surian, J. Leask, A. Dey, K. D. Mandl, and E. Coiera. Mapping information exposure on social media to explain differences in HPV vaccine coverage in the United States. *Vaccine*, 35:3033–3040, 2017.

[76] D. Dylewsky, M. Anand, and C. T. Bauch. Early warning signals for bifurcations embedded in high dimensions. *Sci. Rep.*, 14:18277, 2024.

[77] D. J. D. Earn, P. Rohani, B. M. Bolker, and B. T. Grenfell. A Simple Model for Complex Dynamical Transitions in Epidemics. *Science*, 287(5453):667–670, 2000.

[78] K. M. Eccles and S. Bertazzon. Applications of geographic information systems in public health: A geospatial approach to analyzing MMR immunization uptake in Alberta. *Can. J. Public Health*, 106:e355–e361, 2015.

[79] P. Erdős and A. Rényi. On random graphs I. *Publ. Math. Debrecen*, 6:290–297, 1959.

[80] I. Farahbakhsh, C. T. Bauch, and M. Anand. Tipping points in coupled human–environment system models: a review. *Earth Syst. Dynam.*, 15:947–967, 2024.

[81] C. P. Farrington, H. J. Whitaker, J. Wallinga, and P. Manfredi. Measures of disassortativeness and their application to directly transmitted infections. *Biom J.*, 51:387–407, 2009.

[82] F. Fenner, D. A. Henderson, I. Arita, Z. Ježek, and I. D. Ladnyi. *Smallpox and its eradication*. World Health Organization, 1998.

[83] G. Fenza, M. Gallo, V. Loia, F. Orciuoli, and E. Herrera-Viedma. Data set quality in machine learning: consistency measure based on group decision making. *Appl. Soft Comput.*, 106:107366, 2021.

[84] A. P. Fiebelkorn, S. B. Redd, P. A. Gastañaduy, N. Clemmons, P. A. Rota, J. S. Rota, W. J. Bellini, and G. S. Wallace. A comparison of postelimination measles epidemiology in the United States, 2009–2014 versus 2001–2008. *J. Pediatric Infect. Dis. Soc.*, 6:40–48, 2015.

[85] P. E. M. Fine and J. A. Clarkson. Individual Versus Public Priorities in the Determination of Optimal Vaccination Policies. *Am. J. Epidemiol.*, 124(6):1012–1020, 1986.

[86] D. K. Flaherty. The vaccine-autism connection: a public health crisis caused by unethical medical practices and fraudulent science. *Ann. Pharmacother.*, 45:1302–1304, 2011.

[87] B. M. Flores, E. Montoya, B. Sakschewski, N. Nascimento, A. Staal, R. A. Betts, C. Levis, D. M. Lapola, A. Esquível-Muelbert, C. Jakovac, C. A. Nobre, R. S. Oliveira, L. S. Borma, D. Nian, N. Boers, S. B. Hecht, H. ter Steege, J. Arieira, I. L. Lucas, E. Berenguer, J. A. Marengo, L. V. Gatti, C. R. C. Mattos, and M. Hirota. Critical transitions in the Amazon forest system. *Nature*, 626:555–564, 2024.

[88] F. Fu, M. Nowak, N. Christakis, and J. H. Fowler. The Evolution of Homophily. *Sci. Rep.*, 2:845, 2012.

[89] F. Fu, D. I. Rosenbloom, L. Wang, and M. A. Nowak. Imitation dynamics of vaccination behaviour on social networks. *Proc. R. Soc. B*, 278:42–49, 2010.

[90] S. Funk, S. Bansal, C. T. Bauch, K. T. D. Eames, W. J. Edmunds, A. P. Galvani, and P. Klepac. Nine challenges in incorporating the dynamics of behaviour in infectious diseases models. *Epidemics*, 10:21–25, 2015.

[91] S. Funk, E. Gilad, C. Watkins, and V. A. A. Jansen. The spread of awareness and its impact on epidemic outbreaks. *Proc. Natl. Acad. Sci. U.S.A.*, 106(16):6872–6877, 2009.

[92] S. Funk and A. A. King. Choices and trade-offs in inference with infectious disease models. *Epidemics*, 30:100383, 2019.

[93] S. Funk, M. Salathé, and V. A. A. Jansen. Modelling the influence of human behaviour on the spread of infectious diseases: a review. *J. R. Soc. Interface*, 7(50):1247–1256, 2010.

[94] A. P. Galvani, T. C. Reluga, and G. B. Chapman. Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum. *Proc. Natl. Acad. Sci. U.S.A.*, 104:5692–5697, 2007.

[95] P. A. Gastañaduy, J. Budd, N. Fisher, S. B. Redd, J. Fletcher, J. Miller, D. J. McFadden, J. Rota, P. A. Rota, C. Hickman, B. Fowler, L. Tatham, G. S. Wallace, S. de Fijter, Amy Parker Fiebelkorn, and M. DiOrio. A Measles Outbreak in an Underimmunized Amish Community in Ohio. *N. Engl. J. Med.*, 375(14):1343–1354, 2016.

[96] P. A. Gastañaduy, S. B. Redd, A. P. Fiebelkorn, J. S. Rota, P. A. Rota, W. J. Bellini, J. F. Seward, and G. S. Wallace. Measles — United States, January 1–May 23, 2014. *MMWR. Morb. Mortal. Wkly. Rep.*, 63(22):496–499, 2014.

[97] P.-Y. Geoffard and T. Philipson. Disease eradication: private versus public vaccination. *Am. Econ. Rev.*, 87:222–230, 1997.

[98] P. Gerretsen, J. Kim, F. Caravaggio, L. Quilty, M. Sanches, S. Wells, E. E. Brown, B. Agic, B. G. Pollock, and A. Graff-Guerrero. Individual determinants of COVID-19 vaccine hesitancy. *PLoS One*, 16:e0258462, 2021.

[99] M. Giuffrè and D. L. Shung. Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. *NPJ Digit. Med.*, 6:186, 2023.

[100] M. Gladwell. *The Tipping Point: How Little Things Can Make a Big Difference*. Little, Brown, 2002.

[101] J. W. Glasser, Z. Feng, S. B. Omer, P. J. Smith, and L. E. Rodewald. The effect of heterogeneity in uptake of the measles, mumps, and rubella vaccine on the potential for outbreaks of measles: a modelling study. *Lancet Infect. Dis.*, 16:599–605, 2016.

[102] A. Glaubitz and F. Fu. Population heterogeneity in vaccine coverage impacts epidemic thresholds and bifurcation dynamics. *Heliyon*, 9(9):e19094, 2023.

[103] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

[104] C. Granell, S. Gómez, and A. Arenas. Dynamical Interplay between Awareness and Epidemic Spreading in Multiplex Networks. *Phys. Rev. Lett.*, 111(12):128701, 2013.

[105] F. M. Guerra, S. Bolotin, G. Lim, J. Heffernan, S. L. Deeks, Y. Li, and N. S. Crowcroft. The basic reproduction number (R_0) of measles: a systematic review. *Lancet Infect. Dis.*, 17(12):e420–e428, 2017.

[106] G. I. Hagstrom and S. A. Levin. Phase transitions and the theory of early warning indicators for critical transitions, 2021.

[107] V. Hall, E. Banerjee, C. Kenyon, A. Strain, J. Griffith, K. Como-Sabetti, J. Heath, L. Bahta, K. Martin, M.a McMahon, D. Johnson, M. Roddy, D. Dunn, and K. Ehresmann. Measles Outbreak — Minnesota April–May 2017. *MMWR. Morb. Mortal. Wkly. Rep.*, 66(27):713–717, 2017.

[108] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, 2016.

[109] Z. He and C. T. Bauch. Effect of homophily on coupled behavior-disease dynamics near a tipping point. *Math. Biosci.*, 376:109264, 2024.

[110] Z. He and C. T. Bauch. Forecasting infectious disease outbreak risks from vaccine sentiments on social media: A data-driven dynamical systems approach. *Math. Biosci. Eng.*, 22:2761–2779, 2025.

[111] D. Helbing. *Quantitative Sociodynamics: Stochastic Methods and Models of Social Interaction Processes*. Springer Berlin Heidelberg, 2010.

[112] D. Helbing and A. Johansson. Cooperation, Norms, and Revolutions: A Unified Game-Theoretical Approach. *PLoS One*, 5(10):e12530, 2010.

[113] D. Helbing and S. Lozano. Phase transitions to cooperation in the prisoner's dilemma. *Phys. Rev. E*, 81(5):057102, 2010.

[114] D. A. Henderson. *Smallpox: the death of a disease: the inside story of eradicating a worldwide killer*. Prometheus Books, 2009.

[115] H. W. Hethcote. *Modeling heterogeneous mixing in infectious disease dynamics*, page 215–238. Cambridge University Press, 1996.

[116] J. Hilton and M. J. Keeling. Incorporating household structure and demography into models of endemic disease. *J. R. Soc. Interface*, 16:20190317, 2019.

[117] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9:1735–1780, 1997.

[118] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[119] I. Holmdahl and C. Buckee. Wrong but useful – what Covid-19 epidemiologic models can and cannot tell us. *N. Engl. J. Med.*, 383:303–305, 2020.

[120] Y. H. Hsieh. Age groups and spread of influenza: implications for vaccination strategy. *BMC Infect. Dis.*, 10:106, 2012.

[121] M. J. Hutchens, J. D. Hmielowski, and M. A. Beam. Reinforcing spirals of political discussion and affective polarization. *Commun. Monogr.*, 86(3):357–376, 2019.

[122] H. Ismail Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller. Transfer learning for time series classification. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 1367–1376. IEEE, 2018.

[123] H. Ismail Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller. Deep learning for time series classification: a review. *Data Min. Knowl. Dis.*, 33:917–963, 2019.

[124] C. Janiesch, P. Zschech, and K. Heinrich. Machine learning and deep learning. *Electron. Mark.*, 31:685–695, 2021.

[125] V. A. A. Jansen, N. Stollenwerk, H. J. Jensen, M. E. Ramsay, W. J. Edmunds, and C. J. Rhodes. Measles Outbreaks in a Population with Declining Vaccine Uptake. *Science*, 301:804, 2003.

[126] P. C. Jentsch, M. Anand, and C. T. Bauch. Spatial correlation as an early warning signal of regime shifts in a multiplex disease-behaviour network. *J. Theor. Biol.*, 448:17–25, 2018.

[127] N. F. Johnson, N. Velásquez, N. J. Restrepo, R. Leahy, N. Gabriel, S. El Oud, M. Zheng, P. Manrique, S. Wuchty, and Y. Lupu. The online competition between pro- and anti-vaccination views. *Nature*, 582:230–233, 2020.

[128] J. Jurczyk, T. Rehberg, A. Eckrot, and I. Morgenstern. Measuring critical transitions in financial markets. *Sci. Rep.*, 7:11564, 2017.

[129] M. Kalivianakis, S. L. Mous, and J. Grasman. Reconstruction of the seasonally varying contact rate for measles. *Math. Biosci.*, 124:225–234, 1994.

[130] M. J. Keeling and P. Rohani. *Modeling Infectious Diseases in Humans and Animals*. Princeton university press, 2008.

[131] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. A*, 115:700–721, 1927.

[132] A. P. Kinzig, P. R. Ehrlich, L. J. Alston, K. Arrow, S. Barrett, T. G. Buchman, G. C. Daily, B. Levin, S. Levin, M. Oppenheimer, E. Ostrom, and D. Saari. Social norms and global environmental challenges: the complex interaction of behaviors, values, and policy. *BioScience*, 63:164–175, 2013.

[133] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman. 1D convolutional neural networks and applications: a survey. *Mech. Syst. Signal Process.*, 151:107398, 2021.

[134] P. Krishnakumar and S. Karlamangla. How stricter vaccine laws spared California from a major measles outbreak, 2019.

[135] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Proceedings of the 26th International Conference on Neural Information Processing Systems*, page 1097–1105. Curran Associates Inc., 2012.

[136] C. Kuehn. A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics. *Phys. D Nonlinear Phenom.*, 240:1020–1035, 2011.

[137] V. Kumar, C. T. Bauch, and S. Bhattacharyya. A game theoretic complex network model to estimate the epidemic threshold under individual vaccination behaviour and adaptive social connections. *Sci. Rep.*, 14:29148, 2024.

[138] Y. A. Kuznetsov. *Elements of Applied Bifurcation Theory*. Springer, 2023.

[139] S. P. Lalley. Lévy processes, stable processes, and subordinators, 2007.

[140] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.

[141] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86:2278–2324, 1998.

[142] A. D. Lee, N. S. Clemmons, M. Patel, and P. A. Gastañaduy. International Importations of Measles Virus into the United States During the Postelimination Era, 2001–2016. *J. Infect. Dis.*, 219(10):1616–1623, 2018.

[143] S. Lee, M. Golinski, and G. Chowell. Modeling Optimal Age-Specific Vaccination Strategies Against Pandemic Influenza. *Bull. Math. Biol.*, 74:958–980, 2012.

[144] K. Lewis, M. Gonzalez, and J. Kaufman. Social selection and peer influence in an online social network. *Proc. Natl. Acad. Sci. U.S.A.*, 109:68–72, 2012.

[145] T. A. Lieu, G. T. Ray, N. P. Klein, C. Chung, and M. Kulldorff. Geographic clusters in underimmunization and vaccine refusal. *Pediatrics*, 135:280–289, 2015.

[146] R. Liu, P. Chen, K. Aihara, and L. Chen. Identifying early-warning signals of critical transitions with strong noise by dynamical network markers. *Sci. Rep.*, 5:17501, 2015.

[147] R. Liu, J. Wu, and H. Zhu. Media/psychological impact on multiple outbreaks of emerging infectious diseases. *Comput. Math. Methods Med.*, 8:612372, 2007.

[148] L. H. Loft, E. A. Pedersen, S. U. Jacobsen, B. Søborg, and J. Bigaard. Using Facebook to increase coverage of HPV vaccination among Danish girls: An assessment of a Danish social media campaign. *Vaccine*, 38:4901–4908, 2020.

[149] W. P. London and J. A. Yorke. Recurrent outbreaks of measles, chickenpox and mumps. I. Seasonal variation in contact rates. *Am. J. Epidemiol.*, 98:453–468, 1973.

[150] M. S. Majumder, E. L. Cohn, S. R. Mekaru, J. E. Huston, and J. S. Brownstein. Substandard Vaccination Compliance and the 2015 Measles Outbreak. *JAMA Pediatr.*, 169(5):494–495, 2015.

[151] E. Malthouse. Confirmation bias and vaccine-related beliefs in the time of COVID-19. *J. Public Health*, 45:523–528, 2023.

[152] P. Manfredi and A. d'Onofrio. *Modeling the Interplay Between Human Behavior and the Spread of Infectious Diseases*. Springer, 2013.

[153] P. Manfredi, P. D. Posta, A. d'Onofrio, E. Salinelli, F. Centrone, C. Meo, and P. Poletti. Optimal vaccination choice, vaccination games, and rational exemption: An appraisal. *Vaccine*, 28:98–109, 2009.

[154] T. Marks and S. Vanderslott. Which countries have mandatory childhood vaccination policies?, 2021.

[155] J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

[156] R. McDonald, P. S. Ruppert, M. Souto, D. E. Johns, K. McKay, N. Bessette, L. X. McNulty, J. E. Crawford, P. Bryant, M. C. Mosquera, S. Frontin, T. Deluna-Evans, D. E. Regenye, E. F. Zaremski, V. J. Landis, B. Sullivan, B. E. Rumpf, J. Doherty, K. Sen, E. Adler, L. DiFedele, S. Ostrowski, C. Compton, E. Rausch-Phung, I. Gelman, B. Montana, D. Blog, B. J. Hutton, and H. A. Zucker. Notes from the Field: Measles Outbreaks from Imported Cases in Orthodox Jewish Communities — New York and New Jersey, 2018–2019. *MMWR. Morb. Mortal. Wkly. Rep.*, 68(19):444–445, 2019.

[157] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a Feather: Homophily in Social Networks. *Annu. Rev. Sociol.*, 27(1):415–555, 2001.

[158] L. Mello. Fat-Tailed Distributions and Levy Processes, 2008.

[159] J. Menard, T. M. Bury, C. T. Bauch, and M. Anand. When conflicts get heated, so does the planet: coupled social-climate dynamics under inequality. *Proc. R. Soc. B Biol. Sci.*, 288(1959):20211357, 2021.

[160] C. S. Meppelink, E. G. Smit, M. L. Fransen, and N. Diviani. "I was right about vaccination": Confirmation bias and health literacy in online health information seeking. *J. Health Commun.*, 24:129–140, 2019.

[161] C. J. E. Metcalf, M. Ferrari, A. L. Graham, and B. T. Grenfell. Understanding herd immunity. *Trends Immunol.*, 36:753–755, 2015.

[162] L. Mitchell and J. V. Ross. A data-driven model for influenza transmission incorporating media effects. *R. Soc. Open Sci.*, 3:160481, 2016.

[163] B. Mønsted and S. Lehmann. Characterizing polarization in online vaccine discourse–A large-scale study. *PLoS One*, 17:e0263746, 2022.

[164] A. J. Morales, J. Borondo, J. C. Losada, and R. M. Benito. Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos*, 25(3):033114, 2015.

[165] D. M. Morens, G. K. Folkers, and A. S. Fauci. The challenge of emerging and re-emerging infectious diseases. *Nature*, 430:242–249, 2004.

[166] J. Mossong, N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga, et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Med.*, 5:e74, 2008.

[167] M. Müller and M. Salathé. Crowdbreaks: Tracking Health Trends Using Public Social Media Data and Crowdsourcing. *Front. Public Health*, 7:81, 2019.

[168] M. Müller, M. Salathé, and P. E. Kummervold. COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on Twitter. *Front. Artif. Intell.*, 6:1023281, 2023.

[169] M. L. Ndeffo Mbah, J. Liu, C. T. Bauch, Y. I. Tekel, J. Medlock, L. A. Meyers, and A. P. Galvani. The Impact of Imitation on Vaccination Behavior in Social Contact Networks. *PLoS Comput. Biol.*, 8(4):e1002469, 2012.

[170] J. P. Nolan. *Univariate Stable Distributions: Models for Heavy Tailed Data*. Springer, 2020.

[171] Government of Canada. Vaccination coverage in Canada, 2024.

[172] J. K. Olive, P. J. Hotez, A. Damania, and M. S. Nolan. The state of the antivaccine movement in the United States: A focused examination of nonmedical exemptions in states and counties. *PLoS Med.*, 15(6):e1002578, 2018.

[173] T. Oraby, V. Thampi, and C. T. Bauch. The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proc. R. Soc. B Biol. Sci.*, 281(1780):20133172, 2014.

[174] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[175] A. D. Pananos, T. M. Bury, C. Wang, J. Schonfeld, S. P. Mohanty, B. Nyhan, M. Salathé, and C. T. Bauch. Critical dynamics in population vaccinating behavior. *Proc. Natl. Acad. Sci. U.S.A.*, 114(52):13762–13767, 2017.

[176] P. L. Panum. *Observations made during the epidemic of measles on the Faroe Islands in the year 1846*. New York: Delta Omega Society, 1940.

[177] M. Patel, A. D. Lee, N. S. Clemmons, S. B. Redd, S. Poser, D. Blog, J. R. Zucker, J. Leung, R. Link-Gelles, H. Pham, R. J. Arciuolo, E. Rausch-Phung, B. Bankamp, P. A. Rota, C. M. Weinbaum, and P. A. Gastañaduy. National Update on Measles Cases and Outbreaks — United States, January 1–October 1, 2019. *Morb. Mortal. Wkly. Rep.*, 68(40):893–896, 2019.

[178] M. K. Patel, L. Dumolard, Y. Nedelec, S. V. Sodha, C. Steulet, M. Gacic-Dobo, K. Kretsinger, J. McFarland, P. A. Rota, and J. L. Goodson. Progress Toward Regional Measles Elimination – Worldwide, 2000–2018. *MMWR Morb. Mortal Wkly. Rep.*, 68:1105–1111, 2019.

[179] C. I. Paules, H. D. Marston, and A. S. Fauci. Measles in 2019 — Going Backward. *N. Engl. J. Med.*, 380(23):2185–2187, 2019.

[180] N. Perra, D. Balcan, B. Gonçalves, and A. Vespignani. Towards a characterization of behavior-disease models. *PLoS One*, 6:e23084, 2011.

[181] A. Perrin and M. Anderson. Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018, 2019.

[182] R. T. Perry, M. Gacic-Dobo, A. Dabbagh, M. N. Mulders, P. M. Strebel, J.-M. Okwo-Bele, P. A. Rota, and J. L. Goodson. Global control and regional elimination of measles, 2000–2012. *MMWR Morb. Mortal Wkly. Rep.*, 63:103–107, 2014.

[183] V. K. Phadke, R. A. Bednarczyk, D. A. Salmon, and S. B. Omer. Association Between Vaccine Refusal and Vaccine-Preventable Diseases in the United States: A Review of Measles and Pertussis. *JAMA*, 315(11):1149–1158, 2016.

[184] P. Pilipiec, I. Samsten, and A. Bota. Surveillance of communicable diseases using social media: a systematic review. *PLoS One*, 18:e0282101, 2023.

[185] A. J. Pollard and E. M. Bijker. A guide to vaccinology: from basic principles to new developments. *Nat. Rev. Immunol.*, 21:83–100, 2021.

[186] Public Health England. Laboratory confirmed cases of measles, mumps and rubella, England: October to December 2016. *Health Protection Report*, 11(8), 2017.

[187] Public Health England. Laboratory-confirmed cases of measles, rubella and mumps, England: July to September 2018. *Health Protection Report*, 12(42), 2018.

[188] N. Rabb, J. Bowers, D. Glick, K. H. Wilson, and D. Yokum. The influence of social norms varies with "others"' groups: Evidence from COVID-19 vaccination intentions. *Proc. Natl. Acad. Sci. U.S.A.*, 119:e2118770119, 2022.

[189] T. C. Reluga. An SIS epidemiology game with two subpopulations. *J. Biol. Dyn.*, 3(5):515–531, 2009.

[190] T. C. Reluga, C. T. Bauch, and A. P. Galvani. Evolving public perceptions and stability in vaccine uptake. *Math. Biosci.*, 204:185–198, 2006.

[191] M. Salathé. Digital epidemiology: what is it, and where is it going? *Life Sci. Soc. Policy*, 14:1, 2018.

[192] M. Salathé, L. Bengtsson, T. J. Bodnar, D. D. Brewer, J. S. Brownstein, C. Buckee, E. M. Campbell, C. Cattuto, S. Khandelwal, P. L. Mabry, and A. Vespignani. Digital epidemiology. *PLoS Comput. Biol.*, 8:e1002616, 2012.

[193] M. Salathé and S. Bonhoeffer. The effect of opinion clustering on disease outbreaks. *J. R. Soc. Interface*, 5(29):1505–1508, 2008.

[194] M. Salathé and S. Khandelwal. Assessing Vaccination Sentiments with Online Social Media: Implications for Infectious Disease Dynamics and Control. *PLoS Comput. Biol.*, 7(10):e1002199, 2011.

[195] T. S. Sathyanarayana Rao and C. Andrade. The MMR vaccine and autism: Sensation, refutation, retraction, and fraud. *Indian J. Psychiatry*, 53(2):95–96, 2011.

[196] M. Scheffer. *Critical Transitions in Nature and Society*. Princeton University Press, 2009.

[197] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

[198] M. Scheffer, S. Carpenter, J. Foley, C. Folke, and B. Walker. Catastrophic shifts in ecosystems. *Nature*, 413:591–596, 2001.

[199] W. See, Y. K. Ng, L. Cui, Y. Kita, S. P.-L. Ooi, V. Lee, D. M. K. Heng, and R. T. Pin Lin. Singapore's efforts to achieve measles elimination in 2018. *Western Pac. Surveill. Response J.*, 12, 2021.

[200] L. F. Shampine, I. Gladwell, and S. Thompson. *Solving ODEs with MATLAB*. Cambridge University Press, 2003.

[201] A. Shenfield and M. Howarth. A novel deep learning model for the detection and identification of rolling element-bearing faults. *Sensors*, 20:5112, 2022.

[202] L. Sherrard, J. Hiebert, J. Cunliffe, L. Mendoza, and J. Cutler. Measles surveillance in Canada: 2015. *Can. Commun. Dis. Rep.*, 42(7):139–145, 2016.

[203] L. Sherrard, J. Hiebert, and S. Squires. Measles surveillance in Canada: Trends for 2014. *Can. Commun. Dis. Rep.*, 41(7):157–168, 2015.

[204] E. Shim, G. B. Chapman, J. P. Townsend, and A. P. Galvani. The influence of altruism on influenza vaccination decisions. *J. R. Soc. Interface*, 9:2234–2243, 2012.

[205] E. Shim, J. J. Grefenstette, S. M. Albert, B. E. Cakouros, and D. S. Burke. A game dynamic model for vaccine skeptics and vaccine believers: Measles as an example. *J. Theor. Biol.*, 295:194–203, 2012.

[206] E. Simons, M. Ferrari, J. Fricks, K. Wannemuehler, A. Anand, A. Burton, and P. Strebel. Assessment of the 2010 global measles mortality reduction goal: results from a model of surveillance data. *The Lancet*, 379:2173–2178, 2012.

[207] A. Smith, J. Yarwood, and D. M. Salisbury. Tracking mothers' attitudes to MMR immunisation 1996–2006. *Vaccine*, 25(20):3996–4002, 2007.

[208] R. V. Solé. *Phase Transitions*. Princeton University Press, 2011.

[209] J. Sooknanan and D. M. G. Comissiong. Trending on social media: integrating social media into infectious disease dynamics. *Bull. Math. Biol.*, 82:86, 2020.

[210] E. Southall, T. S. Brett, M. J. Tildesley, and L. Dyson. Early warning signals of infectious disease transitions: a review. *J. R. Soc. Interface*, 18(182):20210555, 2021.

[211] Á. Stefkovics, F. Albert, A. S. Ligeti, B. Dávid, S. Rudas, and J. Koltai. Vaccination homophily in ego contact networks during the COVID-19 pandemic. *Sci. Rep.*, 14:15515, 2024.

[212] M. A. Strassburg. The global eradication of smallpox. *Am. J. Infect. Control*, 10:53–59, 1982.

[213] S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, 2015.

[214] L. H. Sun. Anti-vaccine activists spark a state's worst measles outbreak in decades, 2017.

[215] A. Tasciotti, F. Urban, F. de Dea, L. Bortolussi, G. Caravagna, and A. D'Onofrio. Deep learning-informed Bayesian model-based analysis to estimate superspreading events in epidemic outbreaks. *IEEE Access*, 12:161375–161400, 2024.

[216] G. Thompson. Measles and MMR statistics, 2009.

[217] C. W. Tidd, L. F. Olsen, and W. M. Schaffer. The case for chaos in childhood epidemics. II. Predicting historical epidemics from mathematical models. *Proc. R. Soc. B*, 254:257–273, 1993.

[218] S. A. Truelove, M. Graham, W. J. Moss, C. J. E. Metcalf, M. J. Ferrari, and J. Lessler. Characterizing the impact of spatial clustering of susceptibility for measles elimination. *Vaccine*, 37:732–741, 2019.

[219] UK Health Security Agency. Cover of vaccination evaluated rapidly (COVER) programme: annual data, 2015.

[220] UK Health Security Agency. Measles notifications and deaths in England and Wales: 1940 to 2023, 2024.

[221] I. A. van de Leemput, M. Wichers, A. O. J. Cramer, D. Borsboom, F. Tuerlinckx, P. Kuppens, E. H. van Nes, W. Viechtbauer, E. J. Giltay, S. H. Aggen, C. Derom, N. Jacobs, K. S. Kendler, H. L. J. van der Maas, M. C. Neale, F. Peeters, E. Thiery, P. Zachar, and M. Scheffer. Critical slowing down as early warning for the onset and termination of depression. *Proc. Natl. Acad. Sci. U.S.A.*, 111:87–92, 2013.

[222] E. Velasco, T. Agheneza, K. Denecke, G. Kirchner, and T. Eckmanns. Social media and internet-based data in global systems for public health surveillance: a systematic review. *Milbank Q.*, 92:7–33, 2014.

[223] F. Verelst, L. Willem, and P. Beutels. Behavioural change models for infectious disease transmission: a systematic review (2010–2015). *J. R. Soc. Interface*, 13(125):20160820, 2016.

[224] A. Vespignani. Modelling dynamical processes in complex socio-technical systems. *Nature Phys.*, 8:32–39, 2012.

[225] K. Viswanath, M. Bekalu, D. Dhawan, R. Pinnamaneni, J. Lang, and R. McLoud. Individual and social determinants of COVID-19 vaccine uptake. *BMC Public Health*, 21:818, 2021.

[226] A. J. Wakefield, S. H. Murch, A. Anthony, J. Linnell, D. M. Casson, M. Malik, M. Berelowitz, A. P. Dhillon, M. A. Thomson, P. Harvey, A. Valentine, S. E. Davies, and J. A. Walker-Smith. RETRACTED: Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children. *Lancet*, 351(9103):637–641, 1998.

[227] X. Wang, D. Jia, S. Gao, C. Xia, X. Li, and Z. Wang. Vaccination behavior by coupling the epidemic spreading with the human decision under the game theory. *Appl. Math. Comput.*, 380:125232, 2020.

[228] Z. Wang, M. A. Andrews, Z.-X. Wu, L. Wang, and C. T. Bauch. Coupled disease–behavior dynamics on complex networks: A review. *Phys. Life Rev.*, 15:1–29, 2015.

[229] Z. Wang, C. T. Bauch, S. Bhattacharyya, A. d'Onofrio, P. Manfredi, M. Perc, N. Perra, M. Salathé, and D. Zhao. Statistical physics of vaccination. *Phys. Rep.*, 664:1–113, 2016.

[230] Z. Wang, W. Yan, and T. Oates. Time series classification from scratch with deep neural networks: a strong baseline. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 1578–1585. IEEE, 2017.

[231] P. Warden. Geodict: a simple Python library/tool for pulling location information from unstructured text, 2010.

[232] P. Waszkiewicz, P. Lewulis, M. Górski, A. Czarnecki, and W. Feleszko. Public Vaccination Reluctance: What Makes Us Change Our Minds? Results of A Longitudinal Cohort Survey. *Vaccines*, 10(7):1081, 2022.

[233] J. S. Weitz, S. W. Park, C. Eksin, and J. Dushoff. Awareness-driven behavior changes can shift the shape of epidemics away from peaks and toward plateaus, shoulders, and oscillations. *Proc. Natl. Acad. Sci. U.S.A.*, 117:32764–32771, 2020.

[234] A. E. Wilson, C. U. Lehmann, S. N. Saleh, J. Hanna, and R. J. Medford. Social media: a new tool for outbreak surveillance. *Antimicrob. Steward. Healthc. Epidemiol.*, 1:e50, 2021.

[235] C. Wissel. A universal law of the characteristic return time near thresholds. *Oecologia*, 65:101–107, 1984.

[236] L. Worden, S. F. Ackley, J. Zipprich, K. Harriman, W. T. A. Enanoria, R. Wannier, and T. C. Porco. Measles transmission during a large outbreak in California. *Epidemics*, 30:100375, 2020.

[237] World Health Organization. Measles vaccines: WHO position paper – April 2017. *Wkly. Epidemiol. Rec.*, 92(17):205–227, 2017.

[238] World Health Organization. Measles Outbreak Toolkit, 2019.

[239] World Health Organization. How do vaccines work?, 2025.

[240] World Health Organization. Measles reported cases and incidence, 2025.

[241] World Health Organization. Measles vaccination coverage, 2025.

[242] Y. Xiao, T. Zhao, and S. Tang. Dynamics of an infectious diseases with media/psychology induced non-smooth incidence. *Math. Biosci. Eng.*, 10:445–461, 2023.

[243] C. Yang, Y. Yang, and Y. Li. Assessing vaccination priorities for different ages and age-specific vaccination strategies of COVID-19 using an SEIR modelling approach. *PLoS One*, 16:e0261236, 2021.

[244] Y. Ye, A. Pandey, C. Bawden, D. M. Sumsuzzman, R. Rajput, A. Shoukat, B. H. Singer, S. M. Moghadas, and A. P. Galvani. Integrating artificial intelligence with mechanistic epidemiological modeling: a scoping review of opportunities and challenges. *Nat. Commun.*, 16:581, 2025.

[245] Q. Yin, Z. Wang, C. Xia, and C. T. Bauch. Impact of co-evolution of negative vaccine-related information, vaccination behavior and epidemic spreading in multilayer networks. *Commun. Nonlinear Sci. Numer. Simul.*, 109:106312, 2022.

[246] J. Zipprich, K. Winter, J. Hacker, D. Xia, J. Watt, and K. Harriman. Measles outbreak — California, December 2014–February 2015. *MMWR. Morb. Mortal. Wkly. Rep.*, 64(6):153–154, 2015.

APPENDICES

Appendix A

Supplementary Material: Effect of homophily on coupled behavior-disease dynamics near a tipping point

A.1 Model derivation

We model the dynamics of vaccinating behaviors by using evolutionary game theory, where individuals can either play a pro-vaccine strategy (P) or an anti-vaccine strategy (A). Individuals are also divided into two groups, namely "social media users" (social group 1) or "social media non-users" (social group 2), based on whether or not they have posted on social media such as Twitter. We let

$$\begin{aligned} x_1 &= \text{proportion of pro-vaccine individuals using social media,} \\ x_2 &= \text{proportion of pro-vaccine individuals not using social media,} \\ y_1 &= \text{proportion of anti-vaccine individuals using social media,} \\ y_2 &= \text{proportion of anti-vaccine individuals not using social media} \end{aligned}$$

with $x_1 + y_1 = x_2 + y_2 = 1$. The group pressure function f_i^S is

$$f_i^P = x_i + (1-h)\sum_{j \neq i} x_j,$$

$$f_i^A = y_i + (1-h)\sum_{j\neq i} y_j,$$

which are simply the sum of the proportions of pro-vaccine (or anti-vaccine) individuals in the population. Due to homophily, the pressure from different groups is decayed by a factor of $(1-h)$. The payoffs are defined as

$$\begin{aligned}
E_1(P) &= -r_v + \delta_0[x_1 + (1-h)x_2],\\
E_1(A) &= -cI(t) + \delta_0[y_1 + (1-h)y_2],\\
E_2(P) &= -r_v + \delta_0[(1-h)x_1 + x_2],\\
E_2(A) &= -cI(t) + \delta_0[(1-h)y_1 + y_2],
\end{aligned}$$

where $E_i(S)$ denotes the payoff of adopting the strategy S in social group i. r_v is the perceived probability of significant morbidity from the vaccine, δ_0 quantifies the effect of injunctive social norms, $I(t)$ is the proportion of infectious people at time t, and c is some constant, indicating that the probability of being infected is proportional to the current infectious population density. $h \in [0,1]$ represents the level of homophily.

Based on the discussion on pair interactions [111], an imitation process can be described by the Master Equation along with transition rates as follows:

$$\frac{d}{dt}P_a(i,t) = \sum_{i'}\left[w^a(i|i';t)P_a(i',t) - w^a(i'|i;t)P_a(i,t)\right], \tag{A.1}$$

$$w^a(i'|i;t) = \nu^a(i',i;t)R^a(i'|i;t), \tag{A.2}$$

$$\nu^a(i',i;t) = \sum_b \nu_{ab}P_b(i',t), \tag{A.3}$$

$$R^a(i'|i;t) = (1-h(1-\delta_{ab}))\max(E_b(i',t) - E_a(i,t), 0), \tag{A.4}$$

where Eq. (A.1) is the Master Equation for strategy i over time in group a, and Eq. (A.2) is the transition rate from strategy i to strategy i' over time in group a. Eq. (A.3) describes the imitation process, where ν_{ab} denotes the imitation rate between an individual from group a and an individual from group b. The readiness to exchange strategy i to i' is shown in Eq. (A.4) [111, 159], where δ_{ab} is the Kronecker delta. The maximum function makes sure that the change of strategy happens only when a higher payoff is obtained. With Eq. (A.3) and (A.4), Eq. (A.2) becomes

$$w^a(i'|i;t) = \sum_b \nu_{ab}P_b(i',t)(1-h(1-\delta_{ab}))\max(E_b(i',t) - E_a(i,t), 0) \tag{A.5}$$

We plug Eq. (A.5) into Eq. (A.1) and get

$$\begin{aligned}
\frac{d}{dt}P_a(i,t) &= \sum_{i'}\Bigg\{\Bigg[\sum_b \nu_{ab}P_b(i,t)(1-h(1-\delta_{ab}))\max(E_b(i,t)-E_a(i',t),0)\Bigg]P_a(i',t) \\
&\quad -\Bigg[\sum_b \nu_{ab}P_b(i',t)(1-h(1-\delta_{ab}))\max(E_b(i',t)-E_a(i,t),0)\Bigg]P_a(i,t)\Bigg\} \\
&= \sum_{i'}\sum_b (1-h(1-\delta_{ab}))\Bigg\{\Bigg[\nu_{ab}P_b(i,t)\max(E_b(i,t)-E_a(i',t),0)\Bigg]P_a(i',t) \\
&\quad -\Bigg[\nu_{ab}P_b(i',t)\max(E_b(i',t)-E_a(i,t),0)\Bigg]P_a(i,t)\Bigg\} \\
&= \sum_{i'}\Bigg\{\Bigg[\nu_{aa}P_a(i,t)\max(E_a(i,t)-E_a(i',t),0)\Bigg]P_a(i',t) \\
&\quad -\Bigg[\nu_{aa}P_a(i',t)\max(E_a(i',t)-E_a(i,t),0)\Bigg]P_a(i,t)\Bigg\} \\
&\quad +(1-h)\sum_{i'}\sum_{b\neq a}\Bigg\{\Bigg[\nu_{ab}P_b(i,t)\max(E_b(i,t)-E_a(i',t),0)\Bigg]P_a(i',t) \\
&\quad -\Bigg[\nu_{ab}P_b(i',t)\max(E_b(i',t)-E_a(i,t),0)\Bigg]P_a(i,t)\Bigg\}.
\end{aligned}$$

By substituting $i = P$, $i' = A$, $P_a(i,t) = x_1$, $P_b(i,t) = x_2$, $P_a(i',t) = y_1$, $P_b(i',t) = y_2$, and $\nu_{aa} = \nu_{ab} = \kappa$, we get

$$\begin{aligned}
\frac{d}{dt}x_1 &= \kappa x_1 y_1[\max(E_1(P)-E_1(A),0)-\max(E_1(A)-E_1(P),0)] \\
&\quad +(1-h)\kappa\Big\{x_2y_1\max(E_2(P)-E_1(A),0)-y_2x_1\max(E_2(A)-E_1(P),0)\Big\} \\
&= \kappa x_1y_1[E_1(P)-E_1(A)] \\
&\quad +(1-h)\kappa\Big\{x_2y_1\max(E_2(P)-E_1(A),0)-y_2x_1\max(E_2(A)-E_1(P),0)\Big\} \\
&= \kappa x_1(1-x_1)[-r_v+cI+\delta_0(2x_1-1+(1-h)(2x_2-1))] \\
&\quad +(1-h)\kappa\Big\{x_2(1-x_1)\max[-r_v+cI+\delta_0(2-h)(x_1+x_2-1),0] \\
&\quad -(1-x_2)x_1\max[r_v-cI-\delta_0(2-h)(x_1+x_2-1),0]\Big\}
\end{aligned}$$

$$
\begin{aligned}
&\equiv \kappa x_1(1-x_1)[-\omega + I + \delta(2x_1 - 1 + (1-h)(2x_2-1))] \\
&\quad + (1-h)\kappa\Big\{x_2(1-x_1)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\
&\quad - (1-x_2)x_1 \max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\}
\end{aligned}
$$

if we further substitute $\omega \equiv r_v/c$, $\delta \equiv \delta_0/c$, $\kappa \equiv \kappa c$, and use the fact $\max(E_1(P) - E_1(A), 0) - \max(E_1(A) - E_1(P), 0) = E_1(P) - E_1(A)$.

A.2 Model simulation

The number of pro-vaccine social media posts T is calculated by $T = mT_p + k\xi$, where $\xi \sim \text{Stable}(\alpha, \beta, c = c(T_p), \mu = \mu(T_p))$, with stability parameter $\alpha \in (1, 2)$ and skewness parameter $\beta \in (0, 1)$. In the simulation, we choose $\alpha = 1.2$ and $\beta = 0.5$. We set $c(T_p)$ to be a sigmoid-like function

$$
c(T_p) = 3 + \frac{1.5}{1 + \exp\left(-\frac{T_p - 50}{20}\right)},
$$

so $c(T_p)$ is bounded between 3 and 4.5. When the stability parameter $\alpha > 1$, the expectation $\mathbb{E}(\xi) = \mu(T_p)$ at each time t. By letting $\mu(T_p) = (1-m)T_p$ and $k = 1$, the expectation $\mathbb{E}(T) = T_p$. We observed from the empirical data that posting Tweets is more of a random event, so we let $m = 0.1$ be small, making ξ the dominant term in determining the value of T.

Since we also observed a large number of days without social media posts, we randomly assign $T(t) = 0$ for some t in our simulation by introducing a "baseline posting probability" p that is dependent on the social media posts obtained on the previous day, $T(t-1)$,

$$
p = \begin{cases} 0.2, & 0 \le T(t-1) < 10 \\ 0.4, & 10 \le T(t-1) < 50 \\ 0.7, & T(t-1) \ge 50 \end{cases}.
$$

In other words, if there were more social media posts on the previous day, then the probability of posting is higher on the current day for people. Another potential approach is based on the concept of Bernoulli trials. The act of posting denotes success ($X = 1$) and not posting denotes failure ($X = 0$). The probability of success (posting) is linked with

T_p. For instance, it could be the following piecewise function:

$$p = \Pr(X = 1) = \begin{cases} 0.2, & 0 \leq T_p < 10 \\ 0.4, & 10 \leq T_p < 50 \\ 0.7, & T_p \geq 50 \end{cases} .$$

Essentially, the number of social media users engaging in posting activities on any given day is a random variable that follows a binomial distribution $\mathcal{B}(N_1, p)$. We can then multiply the number of individuals who post by a factor denoting the average posts each individual posts to get the estimated number of pro-vaccine posts each day.

The stochastic model $d\mathbf{X} = \mathbf{f}(\mathbf{X}, t)dt + \boldsymbol{\sigma} dL$ was simulated by a self-developed Monte-Carlo algorithm. A discrete version of it can be written as $\mathbf{X}(t+\Delta t) \approx \mathbf{X}(t)+\mathbf{f}(\mathbf{X}(t), t)\Delta t+ \boldsymbol{\sigma}\Delta L(\Delta t)$, where $\Delta L(\Delta t) \equiv L(t+\Delta t) - L(t)$ is the increment of the α-stable Lévy process. Due to the fact that an α-stable Lévy process has stationary increments, $L(t + \Delta t) - L(t)$ has the same distribution as $L(\Delta t)$, which is a random variable that satisfies $L(\Delta t) \sim (\Delta t)^{1/\alpha}\,\text{Stable}(0 < \alpha < 2, \beta = 0)$ [139, 170]. We used the *levy_stable* function in the Python library *scipy.stats* to generate random numbers under a stable distribution with $\alpha = 1.2$, $\beta = 0$. We finally used a clip function to make sure that all solutions lie in the physical space $[0, 1]$.

A.3 Sensitivity analysis

We plot two crucial time points: (i) the first time when the population's pro-vaccine percentage x falls below 95%, and (ii) the first time when the infectious percentage I increases above 10^{-4} when varying δ (social norm strength) and κ (sampling rate) at four different homophily levels (Figure A.1).

To analyze the sensitivity of the two crucial time points to the parameters δ and κ, we compute the numerical elasticity coefficients as $(\Delta Y/Y)/(\Delta X_i/X_i)$, where Y represents the model output, corresponding to one of the crucial time points, and the parameter $X_i \in \delta, \kappa$ (Figure A.2). This metric indicates the percentage change in the output when the parameter is changed by 1%. Notably, we observe that both crucial time points exhibit low sensitivity to κ, with elasticity coefficients falling within the range of $[-1, 1]$ (Figure A.2B and D). Similarly, both time points are less sensitive to small values of δ (Figure A.2A and C).

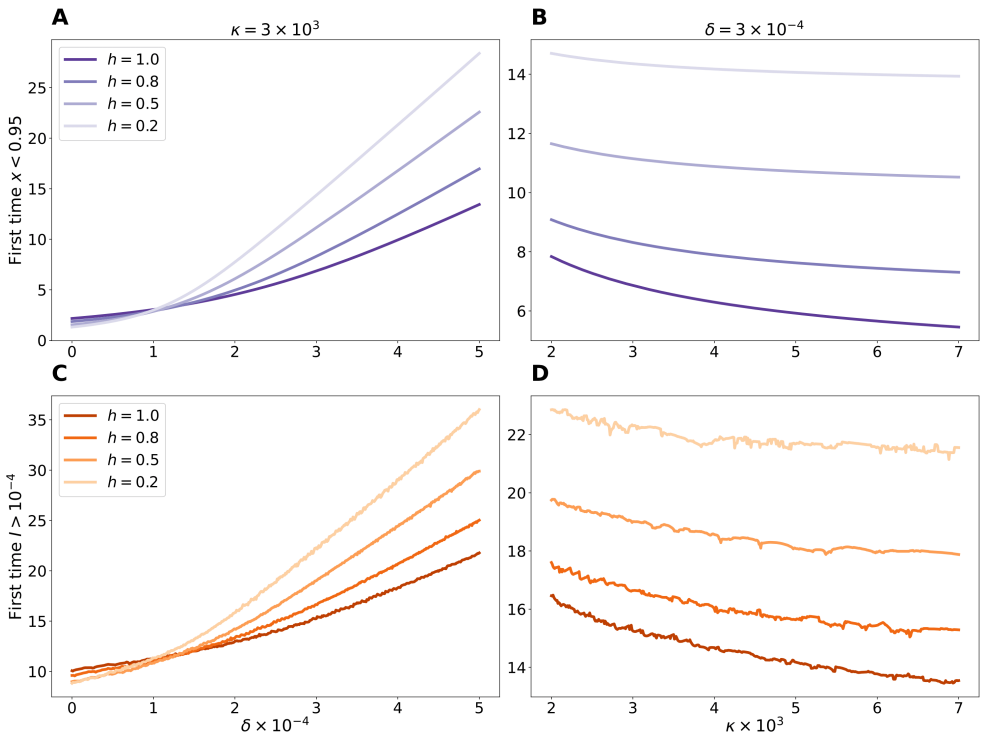


Figure A.1: Simulations of (A, B) the first time when x drops below 95%, and (C, D) the first time when I increases above 10^{-4}, with different values of (A, C) social norm strength δ and (B, D) sampling rate κ.

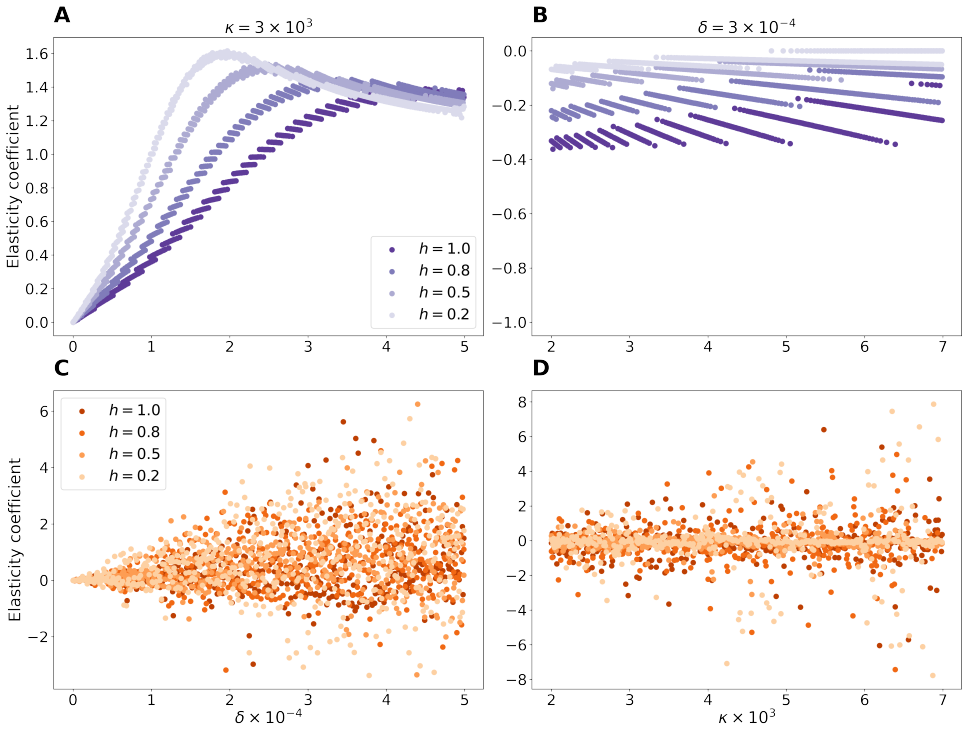


Figure A.2: Corresponding elasticity coefficients of Figure A.1.

Appendix B

Supplementary Material: Homophily in a multi-group vaccination game and its impact on disease dynamics

B.1 Derivation of the game dynamical equations

We describe behavioral changes using the Master Equation [111, 109]:

$$\frac{d}{dt}P_a(i,t) = \sum_{i'}\Big[w^a(i|i';t)P_a(i',t) - w^a(i'|i;t)P_a(i,t)\Big], \tag{B.1}$$

where $P_a(i,t)$ denotes the proportion of individuals in group a who adopt strategy i at time t. The transition rate $w^a(i|i';t)$ captures both spontaneous transitions and imitation processes, and is defined as

$$w^a(i|i';t) = w_a(i|i';t) + \sum_b \nu_{ab}P_b(i,t)R_{ab}(i|i';t), \tag{B.2}$$

where the first term represents the spontaneous strategy changes, and the second term accounts for imitation-based transitions. ν_{ab} is the imitation rate between individuals from groups a and b, typically assumed to be proportional to the relative frequency N_b/N. Unlike [111], we allow for cross-group imitation and define $R_{ab}(i|i';t)$ as the readiness to switch from strategy i to i', based on the positive payoff difference and influenced by the

homophily level h_a:

$$R_{ab}(i|i';t) = \left(\frac{1-h_a}{n} + h_a\delta_{ab}\right)\max[E_b(i,t) - E_a(i',t), 0]. \tag{B.3}$$

$E_b(i,t)$ and $E_a(i',t)$ are expected successes of different strategies. The Kronecker delta δ_{ab} ensures the effect of homophily by amplifying within-group imitation ($a = b$) while reducing the influence of cross-group imitation ($a \neq b$). Plugging Eqs. (B.2) and (B.3) into Eq. (B.1) gives

$$\begin{aligned}
\frac{d}{dt}P_a(i,t) = &\sum_{i'}\Bigg\{\left[w_a(i|i';t) + \sum_b \nu_{ab}P_b(i,t)R_{ab}(i|i';t)\right]P_a(i',t) \\
&- \left[w_a(i'|i;t) + \sum_b \nu_{ab}P_b(i',t)R_{ab}(i'|i;t)\right]P_a(i,t)\Bigg\} \\
= &\sum_{i'}\Bigg\{\left[\sum_b \nu_{ab}P_b(i,t)\left(\frac{1-h_a}{n} + h_a\delta_{ab}\right)\max[E_b(i,t) - E_a(i',t), 0]\right]P_a(i',t) \\
&- \left[\sum_b \nu_{ab}P_b(i',t)\left(\frac{1-h_a}{n} + h_a\delta_{ab}\right)\max[E_b(i',t) - E_a(i,t), 0]\right]P_a(i,t)\Bigg\} \\
&+ \sum_{i'}\Big[w_a(i|i';t)P_a(i',t) - w_a(i'|i;t)P_a(i,t)\Big].
\end{aligned}$$

Since we only consider two strategies, we let $i = P$ (pro-vaccine) and $i' = A$ (anti-vaccine). We further substitute $P_a(i,t) = x_a$, $P_a(i',t) = 1 - x_a$, $P_b(i,t) = x_b$, $P_b(i',t) = 1 - x_b$, $E_a(i,t) = \pi_a(P)$, $E_a(i',t) = \pi_a(A)$, $E_b(i,t) = \pi_b(P)$, $E_b(i',t) = \pi_b(A)$, and $\nu_{ab} = \kappa N_b/N$. We assume that $w_a(i|i';t) = \lambda_a(t)$ is simply a constant or a periodic function, while $w_a(i'|i;t) = \epsilon_a$ is some constant. Therefore, Eq. (B.1) becomes

$$\begin{aligned}
\frac{dx_a}{dt} = &(1-x_a)\sum_b \frac{\kappa N_b}{N}x_b\left(\frac{1-h_a}{n} + h_a\delta_{ab}\right)\max[\pi_b(P) - \pi_a(A), 0] \\
&- x_a\sum_b \frac{\kappa N_b}{N}(1-x_b)\left(\frac{1-h_a}{n} + h_a\delta_{ab}\right)\max[\pi_b(A) - \pi_a(P), 0] \\
&+ \lambda_a(t)(1-x_a) - \epsilon_a x_a \\
= &\left(\frac{1+(n-1)h_a}{n}\right)\frac{\kappa N_a}{N}x_a(1-x_a)\Big\{\max[\pi_a(P) - \pi_a(A), 0] - \max[\pi_a(A) - \pi_a(P), 0]\Big\} \\
&+ \sum_{b\neq a}\left(\frac{1-h_a}{n}\right)\frac{\kappa N_b}{N}\Big\{x_b(1-x_a)\max[\pi_b(P) - \pi_a(A), 0]
\end{aligned}$$

$$
\begin{aligned}
&- x_a(1-x_b)\max[\pi_b(A)-\pi_a(P),0]\Big\} \\
&+ \lambda_a(t)(1-x_a) - \epsilon_a x_a \\
&= \left(\frac{1+(n-1)h_a}{n}\right)\frac{\kappa N_a}{N}\Big\{x_a(1-x_a)[\pi_a(P)-\pi_a(A)]\Big\} \\
&+ \sum_{b\neq a}\left(\frac{1-h_a}{n}\right)\frac{\kappa N_b}{N}\Big\{x_b(1-x_a)\max[\pi_b(P)-\pi_a(A),0] \\
&- x_a(1-x_b)\max[\pi_b(A)-\pi_a(P),0]\Big\} \\
&+ \lambda_a(t)(1-x_a) - \epsilon_a x_a,
\end{aligned}
$$

since $\max[\pi_a(P)-\pi_a(A),0]-\max[\pi_a(A)-\pi_a(P),0]=\pi_a(P)-\pi_a(A)$. Swapping a and b with i and j, respectively, yields the desired equation.

B.2 Matrix representation

We derive a matrix representation of the imitation dynamics, which provides a concise expression and improves computational efficiency when simulating the system, particularly when the number of groups is large. We define the vector of pro-vaccine proportions as $\mathbf{x}=[x_1,x_2,\ldots,x_n]$ and the diagonal pro-vaccine matrix $X=\mathrm{diag}(\mathbf{x})$. The diagonal anti-vaccine matrix is then I_n-X, where I_n is the $n\times n$ identity matrix. We also define the vector of subpopulations $\mathbf{N}=[N_1,N_2,\ldots,N_n]$, and introduce the following matrices:

(i) Contact matrix:

$$
K=\frac{\kappa}{N}\mathbf{1}_n\mathbf{N}=\frac{\kappa}{N}\begin{bmatrix}N_1 & \cdots & N_n\\ \vdots & \ddots & \vdots\\ N_1 & \cdots & N_n\end{bmatrix}.
$$

(ii) Homophily matrix:

$$
H=\frac{1}{n}\begin{bmatrix}1+(n-1)h_1 & 1-h_1 & \cdots & 1-h_1\\ 1-h_2 & 1+(n-1)h_2 & \cdots & 1-h_2\\ \vdots & \vdots & \ddots & \vdots\\ 1-h_n & 1-h_n & \cdots & 1+(n-1)h_n\end{bmatrix}.
$$

(iii) Payoff difference matrix for adopting pro-vaccine strategies:

$$E_{P|A} = \begin{bmatrix} \max[\pi_1(P) - \pi_1(A), 0] & \cdots & \max[\pi_n(P) - \pi_1(A), 0] \\ \vdots & \ddots & \vdots \\ \max[\pi_1(P) - \pi_n(A), 0] & \cdots & \max[\pi_n(P) - \pi_n(A), 0] \end{bmatrix}.$$

(iv) Payoff difference matrix for adopting anti-vaccine strategies:

$$E_{A|P} = \begin{bmatrix} \max[\pi_1(A) - \pi_1(P), 0] & \cdots & \max[\pi_n(A) - \pi_1(P), 0] \\ \vdots & \ddots & \vdots \\ \max[\pi_1(A) - \pi_n(P), 0] & \cdots & \max[\pi_n(A) - \pi_n(P), 0] \end{bmatrix}.$$

Using these matrices, the imitation matrix is given by

$$M = K \odot H \odot [(I_n - X)E_{P|A}X - XE_{A|P}(I_n - X)],$$

in which $\odot$ denotes element-wise (Hadamard) product. Finally, the time evolution of the vector $\mathbf{x}$ is given by

$$\frac{d\mathbf{x}}{dt} = M\mathbf{1}_n,$$

where $\mathbf{1}_n = [1, 1, \ldots, 1] \in \mathbb{R}^n$ is a column vector of ones.

B.3 Numerical methods

When dividing the population into three groups, the rate of change for x_1 is

$$\begin{aligned} \frac{dx_1}{dt} = &\left(\frac{1+2h_1}{3}\right)\frac{\kappa N_1}{N}\Big\{x_1(1-x_1)[\pi_1(P) - \pi_1(A)]\Big\} \\ &+ \left(\frac{1-h_1}{3}\right)\frac{\kappa N_2}{N}\Big\{x_2(1-x_1)\max\left[\pi_2(P) - \pi_1(A), 0\right] \\ &- (1-x_2)x_1\max\left[\pi_2(A) - \pi_1(P), 0\right]\Big\} \\ &+ \left(\frac{1-h_1}{3}\right)\frac{\kappa N_3}{N}\Big\{x_3(1-x_1)\max\left[\pi_3(P) - \pi_1(A), 0\right] \\ &- (1-x_3)x_1\max\left[\pi_3(A) - \pi_1(P), 0\right]\Big\} \end{aligned}$$

and the rates of change for x_2 and x_3 follow a similar form.

The coupled model is simulated using the *ode45* function [200] in MATLAB R2024b. The initial condition is set to $[S(0), I(0), V(0)] = [0.01, 10^{-6}, 0.8]$ when the disease is present. We also consider the important case where the target disease is absent: $I(0) = 0$, which is more challenging for the public health system. The initial pro-vaccine proportions in the pro- and anti-vaccine groups are fixed at 0.95 and 0.2, respectively. For the general group, the initial pro-vaccine proportion is chosen around 0.6875, corresponding to the critical threshold $x^* = (1 + \omega/\delta)/2$ in a homogeneous population [173]. This threshold represents the proportion of pro-vaccine individuals at an equilibrium with partial vaccine coverage and partial susceptibility. We simulate the model over 100 years to observe the long-term behavior of each compartment, even though such a duration exceeds typical time scales for infectious diseases.

As far as the setting of the numerical algorithms is concerned, relative and absolute tolerances are set to 10^{-12} and 10^{-14}, respectively. Since the *ode45* solver adapts its step size dynamically based on error estimates in order to perform well even on stiff dynamical systems, we interpolate the solution onto a uniform time grid with 10000 points at the end of each simulation to make plotting easier.