

Viewing Trends in Graph Connectivity as Early Warnings of Epidemics and Vaccine Crises

by

Brendon C. Phillips

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, 2020

© Brendon C. Phillips 2020

Examining Committee Membership

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

Supervisor(s): Chris T. Bauch
Professor, Department of Applied Mathematics,
University of Waterloo

Internal Member: Francis Poulin
Professor, Department of Applied Mathematics,
University of Waterloo

Mohammad Kohandel
Associate Professor, Department of Applied Mathematics,
University of Waterloo

Internal-External Member: Kumaraswamy Ponnambalam
Professor, Department of Systems Engineering,
University of Waterloo

External Examiner: Julien Arino
Professor, Department of Mathematics,
University of Manitoba

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

Brendon Phillips was the sole author of Chs. 2, 3 and 4, written under the supervision of Dr. Chris T. Bauch with input from Dr. Madhur Anand (University of Guelph).

Material presented in Ch. 2:

This research [445] was completed by Brendon C. Phillips (BCP) with co-investigators Drs. Madhur Anand (MA, University of Guelph) and Chris T. Bauch (CTB, University of Waterloo). CTB conceived the idea and supervised the project. BCP designed the methodology and designed and performed all simulations and data analyses. CTB, MA and BCP interpreted and discussed the results of the work. BP and CTB wrote and edited the manuscript with input from MA. Both CTB and MA provided valuable critical feedback.

Citation: Phillips, B., Anand, M. & Bauch, C.T. Spatial early warning signals of social and epidemiological tipping points in a coupled behaviour-disease network. *Sci Rep* 10, 7611 (2020).

Material presented in Ch. 3:

This research [446] was completed by Brendon C. Phillips (BCP) with substantial contributions from co-investigator Dr. Chris T. Bauch (CTB, University of Waterloo). CTB and BP conceived of the idea and CTB supervised the project. BP designed the methodology, and designed and performed all simulations and data analyses. CTB, MA and BP interpreted and discussed the results of the work. BP and CTB and CTB wrote and edited the manuscript with input from MA. Both CTB and MA provided critical feedback.

Citation: Echo chambers as early warning signals of widespread vaccine refusal in social-epidemiological networks Brendon Phillips, Chris T. Bauch. *medRxiv* 2020.10.17.20214312.

Material presented in Ch. 5:

This research [447] was completed by Brendon C. Phillips (BCP) with substantial contributions from co-investigators Drs. Dillon Browne (DB, University of Waterloo), Madhur Anand (MA, University of Guelph) and Chris T. Bauch (CTB, University of Waterloo). DB conceived the idea and supervised the project along with CTB. BCP and CTB designed the methodology and BCP performed all simulations and data analyses. All authors interpreted and discussed the results of the work. BCP, DB and CTB edited the manuscript with input from MA.

Citation: Model-based projections for COVID-19 outbreak size and student-days lost to closure in Ontario childcare centres and primary schools. Brendon Phillips, Dillon Browne, Madhur Anand, Chris Bauch. *medRxiv* 2020.08.07.20170407.

Abstract

When measles was rampant, suffering apparent, and relief desired, the prospect of vaccination was received with open arms by a grateful public. But it worked *too* well, and opinions slowly diverged; scientists saw aggregate health as proof of the efficacy of intervention, while some of the lay public wondered “*But do we really need this vaccine, though? I don’t see sick people...*” Spurious 1998 research linking the MMR vaccine to autism was published and our dreams of eradication evaporated; the diseases were back to stay. The spread of vaccine disinformation through social networks is immediately apparent and easily exploited, even more so due to the strong assortativity of social networks (both online and face-to-face). Therein lies the focus of this thesis; we investigate different measures of spatial grouping as early warnings signals (EWS) of epidemics through the simulation of social and contact networks and the use of various statistical and graph theoretical tools. Using an agent-based model coupling a binary voting dynamic with an $SIRV_p$ model of infection, we simulate a vaccine preventable disease. Each week, agents are given the opportunity to change opinion to that of a friend, while having potentially disease-spreading interactions with many people. The first study confirms that changes in trend of the Moran’s I, Geary’s C and mutual information statistics give early warnings of the critical transitions representing both vaccine crises and epidemics. This is independent of the strength of an injunctive social norm, though through change point testing we confirm that these warnings come closer to vaccine crises as the norm becomes stronger. We find also that the observable distance between vaccine crisis and epidemic spread decreases as the norm strengthens. Confirmation of these results for other different models boosts our confidence in our results. Our second study shows that graph theoretical changes in incidences of opinion-based communities and echo chambers coincide reliably with outbreaks. Clustering, network modularity and the rate of opinion change also provide EWS of both vaccine crises and epidemics in the population. Due to the immense size and traffic of current social networks, only portions of interactions can be observed at any one time, and therefore our third study tests previously effective signals against an incorporation of vaccine hesitance and network sampling. We find that these identified tools remain good EWS, though experiencing penalties on effectiveness dependent on the sampling rate of the population. In sum, our work provides easily employable tools to predict important negative epidemiological events using readily available data, the best-performing of which is the entropy-based mutual information statistic. Given current and expected events, we believe that this thesis makes a solid contribution to the sparse EWS literature for coupled disease-behaviour systems, as well as providing tools that can be used to inform policy decisions surrounding the mitigation of human folly and critical infection events.

Acknowledgements

I sincerely thank my supervisor Dr. Chris T. Bauch for his guidance, insight and boundless patience over these five years. (Also, cheers to Chris for not laughing during my comprehensive when I blanked on a basic question I answered perfectly the day before. Seriously, there had better be no more exams after this... *ever*.) I'd also like to thank Dr. Madhur Anand for her input on my projects, and Dr. Dillon Browne for the opportunity to participate in an exciting interdisciplinary project.

Thank you to Drs. Francis Poulin, Mohammad Kohandel, Kumaraswamy Ponnambalam and Julien Arino for sitting on my advisory and examination committees. Thanks also to Dr. Kirsten Morris for the opportunity to co-ordinate and teach the course *Mathematics and Music*. It was a joy to teach topics I care deeply about, though many of my students have since described me as helpful and enthusiastic yet "very disorganised" and "hates 8am lectures more than I do". At least I've apologised to eight of them...

To my lab-mates and friends - including Lindsey, Kat, Tom, Paul, Kevin and Natalie, Maliha, Vivek, Keenan and Laura, Jon, John, Vivek, Saptarshi and Christian. Thanks for the office chats, pub nights, research discussions, Sporcle sessions and evenings at Kinkaku and Lancaster. At Western, thanks to Dinesh, Cam, Tyson, Nadeana, Carolyn, Marina and Dr. Chris Kapulkin for your warmth and friendship, and to Master Zivorad Petkovic for your training and mentorship. They've stayed with me over the years.

The Orchestra at UWaterloo and Instrumental Chamber Ensembles helped keep me sane and balanced every semester. Special thanks to conductors Erna van Daele and Dan Warren, ensemble coaches Ben Bolt-Martin and Dr. Laurel Swinden, and flute teacher extraordinaire Margaret Austin; you've helped me to grow as a musician and as a person. Thanks also to the members of the UW Karate & Jujitsu Club, especially Eugene, Alison, Devin, Crystal, Joy, Christian, David, Corin, Kathryn, Bryan and all the Bens (or, whichever Ben reads this).

To Mrs. Tessa King-Inniss at UWI Cave Hill who told me "C'mon man, you can't let it win" in 2011 when my C code segfaulted for the 10th time during that tutorial, know that that one statement has kept me going all these years. Dr. Hussein Thompson, thanks for showing me both the good and evil sides of C++.

Most of all, I'd like to thank my mom Eulene Holder for her money, support, patience, encouragement, frequent health checks and other typical forms of motherly love/mild-to-moderate irritation.

Dedication

To Mom, and her favourite two soul-crushing questions: “When are you going to finish?” and “Weren’t you supposed to be done by now/⟨insert month here⟩??”

Please never change.

Table of Contents

List of Figures	xii
List of Tables	xix
List of Acronyms	xx
List of Symbols	xxi
List of Models used	xxv
1 Introduction	1
1.1 Canonical approaches in infectious disease modelling	2
1.1.1 Compartmental models: 1760 – ????	2
1.1.2 Emergent phenomena and agent-based modelling	5
1.1.3 How about a well-mixed ABM?	5
1.1.4 Behaviour and graph theory	8
1.2 Critical transition and the “epidemic”	9
1.3 A coupled behaviour-infection model	13
1.4 Equilibrium and stopping criteria	18
1.5 Change point testing	19
1.6 Objectives	20
1.7 Outline	20

2	Spatial aggregation can indicate regime shifts	22
2.1	Introduction	23
2.2	Methods	25
2.2.1	Early warning signals	26
2.2.2	Parametrisation	27
2.3	Results	28
2.3.1	Population vaccine immunity status can differ from aggregate vaccine opinion	29
2.3.2	EWS trends identify approaching transitions in both social and infection layers	31
2.3.3	Stronger social norms result in decreased lead distance for all EWS	33
2.3.4	There's no single best EWS per σ value	36
2.3.5	EWS can provide better forewarning than trends in model variables	36
2.4	Discussion	39
3	Community structure predicts social shifts	42
3.1	Introduction	42
3.2	Methods	45
3.2.1	The model	45
3.2.2	Early warning signals	46
3.2.3	Parameters and time series	49
3.3	Results	50
3.3.1	Group size and census predict the social transition	52
3.3.2	Clustering and uncertainty also provide early warnings	55
3.3.3	Finding the best and the worst EWS	57
3.4	Discussion	61
4	Spatial EWS are robust to delay and network sampling	63
4.1	Introduction	63
4.2	Methods	65

4.2.1	Model V4	66
4.2.2	Parametrisation	69
4.2.3	Reconsideration of critical point estimates and lead distance	70
4.2.4	Early warning signals: new and old	74
4.3	Results	77
4.3.1	Non-connectivity-based EWS retain their efficacy	77
4.3.2	Group formation remains a strong EWS	80
4.3.3	Clustering and distance metrics predict the transitions K_*	88
4.3.4	Mutual information, crowned yet again	91
4.4	Discussion	94
5	Model-based projections for COVID-19 outbreak size and class closures in Ontario	98
5.1	Introduction	99
5.2	The model	100
5.2.1	Overview	100
5.2.2	Materials and methods	104
5.3	Results	111
5.3.1	Initial stages of the outbreak	111
5.3.2	Outbreak duration	115
5.3.3	Outbreak size and classroom closure	116
5.3.4	Primary school settings	118
5.3.5	Sensitivity Analysis	120
5.4	Discussion	121
6	Conclusion	124
	References	129
	Appendices	177

A	Spatial aggregation can indicate regime shifts	178
A.1	\mathcal{I} (Moran's I) and \mathcal{C} (Geary's C) in terms of join counts	179
A.2	Comparisons of models V1, V2 and V3	182
A.2.1	Model parameters	182
A.2.2	Intertransition distance $K_p - K_s$	188
A.2.3	Multiple definable transitions K_s and K_p	192
A.2.4	Change point testing and warnings of the EWS	196
A.2.5	Reversibility of the κ -series of the EWS	202
A.2.6	Further comparisons of the EWS' performance	205
A.2.7	EWS performance per σ value	208
A.2.8	A grand comparison	213
A.3	Summation	214
B	Community structure predicts social shifts	215
B.1	Lead distance plots	215
B.2	EWS performance grid plots	223
B.3	EWS performance bar charts	228
C	Spatial EWS are robust to delay and network sampling	233
D	Model-based projections for COVID-19 outbreak size	240
D.1	Sensitivity Analysis	240
D.1.1	Varying α_0 and B_H	240
D.1.2	Varying α_0 and R_{init}	240
D.1.3	Varying α_0 and λ_i	243
	Glossary	248

List of Figures

1.1	Schematics of four different compartmental disease models.	4
1.2	Example of the contact network created by an ABM.	6
1.3	Representation of the graph structure of an ABM featuring random mixing.	7
1.4	Comparing the epidemic threshold and the phase transitions of water.	11
1.5	A representation of epidemic curves.	12
1.6	A plot of COVID-19 cases per day (Dec. 2019) for Canada and Kenya.	13
1.7	Diagrams showing the physical (a) and social (b) dynamics of model V1.	15
1.8	Implementation of the model dynamics used for each stochastic realisation.	17
2.1	Contour plots of the parameter plane for model V2.	29
2.2	Time series demonstrating high sensitivity of the social dynamics of model V2.	30
2.3	Trends of the equilibrium values of some EWS approaching the transitions of the social and infection dynamics of model V2.	31
2.4	Trends of some of the EWS' equilibrium values approaching transitions in the social and infection dynamics of model V2.	32
2.5	Figure showing the decreasing trend of the intertransition gap of model V2.	34
2.6	Mutual information gives the largest lead distance of all EWS for 50% of σ values.	37
2.7	Graph of the trends of χ_{\min}^{Lan} (blue) and χ_{\max}^{Lan} (red) with respect to the value of the social norm σ	38
3.1	Community structure of pro-vaccine agents on a small network.	47
3.2	Example time series of social and infection dynamics of model V1 for social norm $\sigma = 0$	50

3.3	Contour plots of the parameter plane of model V2.	51
3.4	Changes in trend in the κ -series of output variables for model V1.	51
3.5	Graph showing the decreasing trend in the intertransition distance of model V1.	52
3.6	Trends of some measures of connectivity of the social network with respect to the perceived risk of vaccination κ	53
3.7	Trends in the lead distance for each EWS' κ -series with respect to the strength of the social norm σ for model V1.	55
3.8	Trends of connectivity measures with respect to perceived vaccine risk κ for model V2.	56
3.9	Trends in the lead distances of the remaining EWS with respect to the strength of the social norm σ	58
3.10	The proportions of social norm σ values for which each EWS gives the biggest and smallest lead distances.	59
3.11	Grid showing the relative performance of each EWS for model V2.	61
4.1	Representation of the opinion dynamics of model V4.	67
4.2	Implementation of the model dynamics used for each stochastic realisation.	68
4.3	Heat map plots showing the ensemble means of model variables with respect to the social norm σ and the perceived vaccine risk κ for model V4.	69
4.4	The total dynamics of models V3 and V4 at equilibrium show bistability about the perceived vaccine risk $\kappa = 0$ for all parameter values tested.	71
4.5	Unsteady trend of the intertransition distance for model V4.	72
4.6	Trends of the model variables of the social and infection dynamics for different social norms.	73
4.7	Estimated locations of the social and infection transitions drift as the sampling proportion β changes.	73
4.8	Diagram showing the analogue of triadic closure occurring in model V4.	76
4.9	Illustration of the tenuous relationship between the join count and the triad census.	77
4.10	Time series demonstrating the sensitivity of system dynamics to small changes in perceived vaccine risk in the absence of a social norm.	78
4.11	Changes in the trends of mutual information, probabilities of infected neighbours and opinion changes predict transitions in model V4.	79

4.12	Lead distances of non-connectivity-based EWS.	80
4.13	Sudden changes in opinion communities strongly signal upcoming transitions in the dynamics of model V4.	82
4.14	Changes in the sizes and numbers of echo chambers in the network warn of the transitions K_s^1 and K_p^1	83
4.15	Trends in the join counts also give clear warnings of approaching social and infection transitions.	84
4.16	Trends of the lead distances of connectivity-based EWS under the SNHT with respect to the social norm σ	85
4.17	Values of the global clustering coefficient, triad census and opinion network diameter with respect to the perceived vaccine risk.	89
4.18	Trends in the lead distances given by other connectivity-based EWS investigated using the SNHT.	90
4.19	Grid comparing the performance of each EWS at each value of the social norm σ under the SNHT.	92
4.20	Side-by-side bar charts comparing the relative performance of each EWS under the SNHT.	93
5.1	Schematic representation of the model population.	101
5.2	Diagram showing the SEPAIR infection progression for COVID-19.	102
5.3	Plot for the calibration of the β^H parameter.	110
5.4	Time series of the number of exposed and infectious agents in each reopening scenario.	112
5.5	Bar chart of the effective reproduction rate.	114
5.6	Chart of the times taken until the first secondary case.	114
5.7	Box plots of the distribution of simulation durations for each scenario.	115
5.8	Trends in the mean proportions of school attendees in each stage of disease progression.	116
5.9	Mean number of infections produced at each model location.	117
5.10	Box plots of the number of student days lost due to classroom outbreak.	117
5.11	Bar chart showing the durations of room closures in the school due to outbreak.	118
5.12	Further effects of varying class size when cohorting is introduced.	119

5.13	Time series of the trends of each stage of disease progression.	120
5.14	Box plots of the distribution of simulation durations for each scenario.	121
A.1	With no social norm, varying noise parameter ξ_{52} does not affect model V3 when the perceived vaccine risk is close to zero.	184
A.2	In the absence of a social norm, increasing infectivity does not alter sensitivity of the social dynamics of model V1, but does change the vaccination rate when κ becomes positive.	185
A.3	Contour plots of a subregion of the parameter space of model V2.	186
A.4	Contour plots of $\langle V_p \rangle$ and $\langle V_s \rangle$ for the models V1 and V3 showing that investigated parameter regions capture the transitions K_s and K_p	187
A.5	Trends of the EWS' equilibrium values in model V1 with infectivity $p = 0.2$ approaching the transitions of the social and physical dynamics.	188
A.6	Demonstration of the trends of the EWS of model V1 with infectivity $p = 0.8$ approaching the transitions of the social and physical dynamics.	189
A.7	Demonstration of the vanishing intertransition distance $K_p - K_s$ for all models.	191
A.8	For model V2, multiple potential transitions K_s and K_p can be identified for $\sigma > 1$, introducing the opportunity for false warnings.	193
A.9	Similar to model V2, multiple social ($\#K_s > 1$) and physical ($\#K_p > 1$) transitions occur in realisations of model V1 ($p = 0.2$) at $\sigma \geq 0.75$, again providing opportunities for false warnings.	194
A.10	Bar chart showing the number of physical ($\#K_p$, red) and social ($\#K_s$, blue) transitions with respect to the value of σ	195
A.11	Trends of the predicted warnings of the EWS with respect to the length of the κ -series used.	197
A.12	The lead distances given by different change point detection tests for model V2.	199
A.13	The lead distances given by different change point detection tests for model V1, with infectivity $p = 0.2$	200
A.14	The lead distances given by different change point detection tests for model V1, with infectivity $p = 0.8$	201
A.15	The skewness of the EWS' trends varies more consistently with the strength of the social norm than with the intertransition distance for model V2.	203

A.16	Trends of the skewness γ_1 of the κ -series with respect to the strength of the social norm σ for models V1 with $p = 0.2$, $p = 0.8$ and V2.	204
A.17	Comparisons of the performance of EWS (WS) and model dynamics (S/I) for model V2 with respect to four change point detection tests.	206
A.18	Similar to Fig. A.17, we show comparisons of the performance of EWS (WS) and model dynamics (S/I) for model V1 with respect to four change point detection tests.	207
A.19	Bar charts showing the ratios of σ values for which each EWS gave the largest (green bars) and smallest (red bars) lead distance of all EWS for model V2.	209
A.20	Bar charts showing the proportions of σ values for which each EWS gives the largest and smallest lead distances of all EWS for model V1, under the four change point tests.	210
A.21	Grid plots showing the relative performance of each EWS with respect to the value of the social norm σ for model V2.	211
A.22	Grid plots showing the relative performance of each EWS with respect to the value of the social norm σ for model V1.	212
B.1	Trends in the lead distances of the sizes of different types of opinion communities under different change point tests.	216
B.2	Trends in the lead distances of the sizes of different types of echo chambers under different change point tests.	217
B.3	Trends in the lead distances of the censuses of different opinion communities under different change point tests.	218
B.4	Trends in the lead distances of the number of different opinion changes under different change point tests.	219
B.5	Trends in the lead distances of the global clustering coefficient under different change point tests.	220
B.6	Trends in the lead distances of the modularity statistic under different change point tests.	221
B.7	Trends in the lead distance given by the probability of having an infected neighbour under different change point tests.	222
B.8	Performance of each EWS per σ value with the SNHT.	224
B.9	Performance of each EWS per σ value with the Lanzante test.	225

B.10	Performance of each EWS per σ value with the Pettitt test.	226
B.11	Performance of each EWS per σ value with the Buishand range test.	227
B.12	Grids showing the absolute and relative performance of the EWS for different infectivities under the SNHT.	229
B.13	Grids showing the absolute and relative performance of the EWS for different infectivities under the Pettitt test.	230
B.14	Grids showing the absolute and relative performance of the EWS for different infectivities under the Lanzante test.	231
B.15	Grids showing the absolute and relative performance of the EWS for different infectivities under the Buishand range test.	232
C.1	Locations of the physical and social transitions (K_p^β and K_s^β), and trends of the intertransition distance $K_p^\beta - K_s^\beta$ obtained by sampling different proportions β of agents in the network.	234
C.2	Changes in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (1).	235
C.3	Changes in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (2).	236
C.4	Variance in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (3).	237
C.5	Performance of the EWS with respect to the value of the social norm σ when 60% of the network is sampled.	238
C.6	Performance of the EWS with respect to the value of the social norm σ when 80% of the network is sampled.	239
D.1	Results of varying the parameters β^H and α_0 on the number of produced infections for RA allocation.	241
D.2	Results of varying the parameters β^H and α_0 on the number of produced infections for ST allocation.	242
D.3	Results of varying the parameters R_{init} and α_0 on the number of infections for ST allocation.	244
D.4	Results of varying the parameters R_{init} and α_0 on the number of infections for RA allocation.	245

D.5	Results of varying the parameters λ_i and α_0 on the number of infections for ST allocation.	246
D.6	Results of varying the parameters λ_i and α_0 on the number of infections for RA allocation.	247

List of Tables

1.1	Convention for naming the compartments in epidemiological models.	3
1.2	Table of all variables appearing in Chs. 2, 3, and 4.	18
1.3	Table of all parameters appearing in Chs. 2, 3 and 4.	19
4.1	Lead distances and sampling penalties for mutual information, the probability of having an infected neighbour and opinion changes for different sampling penalties.	81
4.2	Lead distances and sampling penalties of the mean sizes of opinion communities and echo chambers for different sampling proportions.	86
4.3	Lead distances and sampling penalties of the numbers of opinion communities and echo chambers for different sampling proportions.	87
4.4	Lead distances and sampling penalties of join count statistics for different sampling proportions.	88
4.5	Lead distances and sampling penalties of the opinion network diameter, triad census and global clustering coefficient for different sampling proportions.	97
5.1	Table showing the various reopening scenarios tested.	105
5.2	Parameter definitions, baseline values and literature sources.	108
5.3	Peak times and values of exposed and infectious agents during the first 30 days of infection spread.	113
A.1	A table of the baseline parameter values used for each simulation.	182
A.2	Table showing the differences of the proportions of maxima and minima for each EWS, with respect to model and change point test.	213

List of Acronyms

ABM Agent-based model.

EBM Equation-based model.

ER Erdős-Rényi network structure.

EWS Early warning signal.

GCC Global clustering coefficient.

GDP Gross domestic product.

SNHT Standard normal homogeneity test.

TVG Ternary voting game.

WHO World Health Organisation.

List of Symbols

$\mathbf{Lan}_\sigma\{\Psi\}$ Lanzante change point test.

σ Strength of an injunctive social norm.

κ The perceived risk of vaccine adverse events.

ξ Probability of switching opinion randomly.

V_s Pro-vaccine opinion.

N Anti-vaccine opinion.

$\mathbb{P}_n(N \rightarrow V_s)$ Probability of switching from anti-vaccine to pro-vaccine opinion.

$\mathbb{P}_n(V_s \rightarrow N)$ Probability of switching from pro-vaccine to anti-vaccine opinion.

n An agent in the population.

α Probability of entering the simulation as a vaccinated pro-vaccine agent (V_s, V_p).

I_n The number of infected neighbours of agent n .

D The size of the network; that is, the number of vertices in the corresponding graph.

d_n^* The number of social contacts of agent n with opinion $*$.

d_n The number of social contacts of agent n .

S Susceptible epidemiological status.

I Symptomatically infectious epidemiological status.

R Recovered epidemiological status.

V_p Vaccinated epidemiological status.

P Presymptomatic epidemiological status.

E Exposed epidemiological status.

A Asymptomatically infectious epidemiological status.

ℓ Duration of the illness.

p The probability of infection. Also called **infectivity**.

μ Probability of birth and death (also termed *replacement*).

ι Proportion of susceptibles infected to mimic case importation.

\mathcal{M} **Mutual information**.

\mathcal{I} **Moran's I**.

W The number of edges in the network.

C **Geary's C**.

K_s Transition point of the social dynamics in the coupled behaviour-infection model defined in Chs. 2 and 3 as the smallest κ value at which $\langle V_s \rangle \approx \langle N \rangle$ (the mean number of pro-vaccine agents equals the number of anti-vaccine agents). Ch. 4 defines K_s as the minimum of all estimates of the transition point across all samples of the network, that is $K_s = \min_{0 < \beta \leq 1} K_s^\beta$.

K_p The transition point of the infection dynamics in the coupled behaviour-infection model defined in Chs. 2 and 3 as the earliest κ at which $\langle R \rangle \approx \langle V_p \rangle$.

$\mathbf{Lan}_\sigma \{ \langle \Psi \rangle \}$ A change point obtained from the Lanzante test.

WS The class of all **early warning signals**.

S/I The collection of all **model variables**.

Q Network **modularity**.

C_* The **global clustering coefficient**.

Θ_* The number of opinion changes performed by agents with opinion $*$.

Γ_* The total probability of an agent with the vaccine opinion $*$ having at least one infected neighbour.

J_* The set of all echo chambers of agents with vaccine opinion $*$.

- Z_* The set of all (opinion-based) communities of agents with vaccine opinion $*$.
- SNHT** $_{\sigma}\{\langle\Psi\rangle\}$ The standard normal homogeneity [change point detection test](#).
- H The hesitant vaccine sentiment. Hesitant agents do not vaccinate, nor do they exert social pressure on their social contacts.
- β The proportion of the social network randomly sampled to compute [early warning signals](#).
- K_s^β The transition point in the social dynamics K_s as obtained from a random sample of $\beta \cdot N$ agents, where N is the size of the network.
- K_p^β The transition point in the infection dynamics K_p as obtained from a random sample of $\beta \cdot N$ agents, with D the network size.
- K_M The earliest transition in the model dynamics, that is, the smaller of K_s and K_p .
- LLT** The proportion of social norm σ values for which an [EWS](#) gives the largest lead distance of all EWS.
- ULT** The proportion of social norm σ values for which an [EWS](#) gives undesirable warnings (that is, undefined, negative or minimum lead distances).
- Pen** The [performance penalty](#) incurred by an [early warning signal](#) upon sampling a proportion β of the population.
- \mathcal{R}_e The [effective reproduction number](#).
- $[X, Y]$ The [join count](#); specifically, the number of friendships between agents with vaccine stances X and Y respectively.
- Ω_* The [graph diameter](#).
- Δ_* The [triad census](#).
- η The probability of becoming symptomatically infected.
- δ The probability of becoming presymptomatically infectious per time step.
- ST** The classroom allocation strategy of rooming [siblings together](#).
- RA** The [random assignment](#) strategy of rooming students.
- β_{ij}^H The probability that an infectious agent of age i will transmit COVID-19 to a flatmate of age j .

- β_{ij}^C The probability that an infectious school attendee of age i will transmit COVID-19 to a classmate of age j .
- β_{ij}^O The probability that an infectious school attendee of age i will transmit COVID-19 to an fellow attendee of age j in a common area.
- λ_S The probability of community infection for a school attendee.
- λ_N The probability of community infection for an agent not attending the school.
- R_{init} The initial proportion of agents assumed recovered due to previous infectious spread in the region.
- o The proportion of childless educators in the institution.
- σ The probability of progression to an infectious epidemiological status P of I per day.
- γ_A The probability of recovery for an asymptotically infected agent per day.
- γ_I The probability of removal (due to mandatory self-isolation) for a symptomatically infected agent per day.
- $\mathbf{Lead}_\sigma[\langle \Psi \rangle]$ The **lead distance** of some **early warning signal** Ψ .

List of Models used

- V1** The basic [ABM](#) outline presented in Sec. 1.3, with size $D = 10000$. This model is used to obtain and test various [EWS](#) in Chs. 2 and 3.
- V2** Identical to model [V1](#), but with size $D = 40000$. This model is used for validation of the [EWS](#) evaluated in App. 2, with corresponding findings given in App. A and auxiliary results in App. B.
- V3** Identical to model [V1](#), but with size $D = 562500$ (that is, 750^2). Similarly to V2, model V3 is also used for validation of the [EWS](#) evaluated in Ch. 2, with corresponding findings also given in App. A.
- V4** A model of size $D = 10000$ featuring an infection dynamic identical to [V1](#), which introduces vaccine hesitance and a small-world network structure to the social dynamics of V1 (fully described in Sec. 4.2). Used in Ch. 4, it also tests the [EWS](#) derived in Chs. 2 and 3 to sample various proportions of randomly chosen agents in the network (100%, 80% and 60%). Auxiliary and confirmatory materials are given in App. C.
- SEPAIR** A well-mixed [ABM](#) model of [COVID-19](#) transmission in childcare centres and primary schools in Ontario. Used in Ch. 5.

Chapter 1

Introduction

Large-scale infectious disease events are extremely detrimental to the region of the outbreak, as well as causing extensive economic and infrastructural damage [355]. The desire to avoid these gross impacts creates the interest in (and need for) tools that can use readily available information to predict their occurrence. This motivates our research into finding *early warning signals* (EWS) of an epidemic. *Warning signals* (both spatial and temporal) are statistically significant, recognisable and characteristic behaviours known to accompany critical transitions in dynamical systems; those preceding the transition are called *early warning signals* [484].

Currently, the world is in the grips of a new plague: most frequently called “Coronavirus” in social and popular media, COVID-19 (the disease caused by the SARS-CoV-2 virus) has taken on a life of its own, causing mass sickness, death and public health crises while also inspiring global panic, paranoia, conspiracy theories, racism, comedy sketches and memes. First identified in December 2019 in Wuhan, a city in the Chinese province of Hubei, the SARS-CoV-2 virus was first thought to have originated in bats (a known source of betacoronaviruses) [494, 19]. Soon, cases of infection were found that could not be traced to any known index cases [329]. Inevitably, the virus soon made its way across Asia, Europe and the Americas, leaving thousands of deaths in its wake. The World Health Organisation (WHO) deemed the outbreak as “a public health emergency of international concern” on 30 January 2020 [10], and publicly declared it a “pandemic” (a disease with global prevalence) on 11 March 2020 [418, 135].

Criticism of the actions of the WHO, pervasive tracking efforts by multiple public health systems and world organisations, large-scale preventative measures and policy interventions [407, 520, 8], intense media scrutiny [375, 581, 214] and data emerging from those countries hardest hit (at the time) by casualty (such as Iran, South Korea, the People’s Republic of China and Italy) [131, 564, 411, 47, 279] offer an unprecedented view into the social impacts, decision making, policy deployment and the gargantuan expense involved in having to employ scalable interventions like population-wide lockdowns, once early efforts at containment have failed.

One aspect of disease control mentioned frequently is vaccination. Currently there are many clinical trials of [vaccines](#) attempting to prevent [COVID-19](#) infection through various avenues [[326](#), [186](#)], many expected to be successful, with companies pre-emptively increasing production capacity to supply the necessary number of doses [[544](#), [538](#), [346](#)], and countries vying to obtain control over distribution in a concept called “vaccine nationalism” [[513](#), [323](#)]. In the absence of a vaccine, physical distancing guidelines have been encouraged and enforced, leading to protests against the perceived wrongs of confinement [[172](#), [233](#)]. One cause of concern directly related to this thesis is the alliance forged by coronavirus protesters and anti-vaccination activists [[32](#), [260](#)].

In what were previously seen as separate causes, there lie some similarities such as conspiracy theories and resistance to measures thought of as excessively restrictive and/or harmful, with some media outlets outlining coronavirus vaccine conspiracies and resistance formed while vaccine hopefuls are still in trial [[292](#), [449](#)]. During the distribution of any resulting vaccine, we can therefore well expect staunch resistance and an increased spread of disinformation. Thereby, we can safely predict that social [network](#) mining and large-scale data gathering will see increased attention and aggressive pursuit, alongside (increased) investment in the development of surveillance and warning systems by government agencies and non-profit organisations alike.

1.1 Canonical approaches in infectious disease modelling

Various studies employ different mathematical frameworks to model and analyse the transmission of infectious disease. A large majority of these leverage knowledge and observation of the disease process to subdivide its progression into stages called *compartments*.

1.1.1 Compartmental models: 1760 – ????

Compartmental modelling of disease spread forms the backbone of current research in mathematical epidemiology. This approach partitions the population into *compartments* representing different stages of the disease process. The first recorded instance of compartmentalisation¹ was done by Daniel Bernoulli in 1766², used to model the effect of vaccination on the spread of smallpox [[60](#), [161](#), [29](#)]. This model broke the population into the compartments S (susceptible) and I

¹the process of modelling where a population is divided into groups of individuals with different stages of progression of the disease.

²This is not directly asserted in [[29](#)], but rather inferred from the language used. The common assertion that this model is first is acknowledged in [[161](#)] and [[225](#)], though not discussed in either article. However, [[225](#)] is careful to state that Bernoulli is not considered to be the first mathematical epidemiologist (in the commonly understood sense); an example would be his contemporary Jean le Rond d’Alembert [[161](#)]. Bernoulli’s research was primarily motivated by the sale of life insurance (annuities) [[60](#), [161](#)], while d’Alembert’s approach is rooted in risk aversion and optimisation [[435](#), [26](#)]. [[225](#)] comments on the decidedly political and economic tone of Bernoulli’s approach, although that can be said of both scientists, who used both economic and moral arguments. The major difference between the two models is Bernoulli’s assumption of the independence of age and the probability of


Compartment	Interpretation
C	a <i>carrier</i> of the disease
E	the agent has been <i>exposed</i> to the pathogen
I	<i>infected with the disease</i> and infectious (can pass the disease to others)
M	has <i>maternally-derived immunity</i> derived from the mother's antibodies
R	<i>recovered</i> from the disease, or <i>removed</i>
S	<i>susceptible</i> to infection
V	<i>vaccinated</i> (immune the disease during the vaccine's effective period)

Table 1.1: Convention for naming the compartments in compartmental epidemiological models. The disease process used is represented by a chain of characters; for example, a model process with susceptible, infected and recovered compartments is labelled ‘SIR’. Examples of such models are shown in Fig. 1.1c


(immune), and treated the relationship between age and susceptibility (for an endemic infection with a small mortality rate) [161]. Since then, many compartmental models have been employed in epidemiological research, many of them centred around threshold effects in diseases spreading through SIR, SIS (Fig. 1.1a), MSIR and SIRS epidemic models on various network structures [291]; the names of the compartments are given in Table 1.1.

As a simple example of model construction, the process of the measles disease is sometimes modelled by an SEIR model [266], where each agent rotates through four distinct stages of the disease (called *compartments*). Some agents are initially *susceptible* to infection (compartment S), while those agents who have been exposed to the disease (infected) pass through two compartments (as shown in Fig. 1.1c): the *exposed* compartment (E) represents the period where the agent is not yet infectious. Agents exit this stage upon becoming infectious, which is sometimes considered roughly coincident with the onset of symptoms (such as a characteristic rash), and is so treated as a characteristic 10-14 *latency period* in simplified models [266, 440], though this is not strictly correct. The agent enters the *infectious* compartment (I) in the second stage of infection (also called the *morbid stage*) when the measles rash appears, where they remain for 5-8 days [440, 266] (this represents a simplification, since a person exposed to the measles pathogen is deemed infectious anywhere from 2 to 4 days *before the onset of symptoms*). The agent then enters the *recovered* compartment (R) after the course of infection [266], where they have gained natural active immunity.


death, in contrast to d’Alembert’s inclusion of marginal utility and time discounting [475, 532]. These early works on probability are given much more prominence in economic theory and the origin of probability theory, than in epidemiology.




(a) SI model, sometimes used to model herpes, or any other disease from which there is no recovery.



(b) SIS model, sometimes used to model gonorrhea [255] and meningitis [483].



(c) SEIR model, sometimes used to model the Ebola [464, 158] and Influenza A [180] diseases.



(d) SEIRC models (such as the one shown above) can be used to model the chikungunya [178] and foot-and-mouth [354] diseases.

Figure 1.1: Schematics of four different compartmental disease models used in mathematical epidemiology. In (c), β represents the rate at susceptible agents are exposed to the infection (*transmission rate* or *force of infection*), σ represents the rate at which latent agents become infected (*incubation rate*), γ the rate at which ill agents recover from illness (*incubation rate*) and ξ represents the rate at which agents lose their immunity and become susceptible once again (*waning rate*).

1.1.2 Emergent phenomena and agent-based modelling

Agent-based models (ABMs) allow us to explore the complex phenomena caused by the interaction of *agents*, which are heterogeneous computational units with individual properties that follow well-defined rules) [136]. Such agents interact with each other and some surrounding environment (constraints, structures and parameter faced by all agents of the simulation) [266, 149, 73] through time-structured agent-agent contacts and community [479, 559, 159, 276]. As such, the complexity of ABMs allow us to capture and measure many deterministic elements of the spreading process at once, such as spread within cohorts, age dependence, the effect of socioeconomic status, as well as identifying target demographics for intervention (looking at shared characteristics) [265, 137, 267].

ABMs have been used to model diseases such as influenza [208, 159, 23, 322], human papillomavirus [438, 227, 252, 92], West Nile virus [75, 331], Ebola [499, 266, 542, 276], plague [320, 560] and others [559], sometimes with the intent of testing strategies for preventing or mitigating outbreaks and the associated problems caused [357, 559, 2, 558, 89, 107, 98]. We use the agent-based simulation paradigm to model heterogeneous mixing among individuals with different opinions and states of health, as well as to impose a random (Erdős-Rényi, ER) network structure on each of the disease and communication networks; not only does this *bottom-up* modelling approach allow the capture of elusive analytically intractable dynamics not captured by differential equation models [466, 430, 193, 73, 157], but also allows us to view the individual trajectory of each agent [215, 472, 73, 107] and track the movement of agents and interactions in dynamic models [157, 28, 433, 169].

Through agent-agent interactions, an ABM will have an inherent connectivity structure among the group of agents (population) represented. Since we are modelling infection spread, we define *effective contact* as any contact between agents that allows for infection of a susceptible individual (such as a cough, sneeze, kiss, shared needle, unprotected sex, interaction with a pet, etc.), depending on the routes of transmission of the disease being modelled. Visually, this is represented by a graph such as Fig. 1.2, where nodes (green squares) represent the agents, and edges (black solid lines) represent some contact between the two agents it connects. An infection process on networks can be cast as *site (or bond) percolation* processes [285]. In site percolation, individuals (agents) are represented by *sites* (squares in Fig. 1.2); a site is considered occupied if the corresponding agent is infected, and empty otherwise. Meanwhile, the disease spreads along the *bonds* (edges in Fig. 1.2) again representing effective contacts made between two agents.

1.1.3 How about a well-mixed ABM?

In general disease modelling, one common simplification is to assume a *well-mixed* population, in which *every* susceptible individual has effective contact with *every* infected agent in (a well-specified subgroup of) the population. This methodology plays a large part in Ch. 5’s model of COVID-19 infection and disease transmission; within the context of agent-based modelling,




Figure 1.2: Example of the contact network created by an [ABM](#). Each square represents an agent (individual) in the model, and each edge represents an effective contact between two agents. Green squares are susceptible to the disease, and red squares are infected with the disease.

it is most often titled *random mixing* or *complete mixing*, as in [213, 104, 534]. In a spatially extended system, any area in which agents are well-mixed presents a *clique*, a group of agents whose induced subgraph is a complete graph K_n (that is, all agents are pairwise connected) with no other agent in the network connected to *all* the agents in the group [295]. Many ABMs feature clique formation through the progression of model dynamics; Sec. 1.2 presents many examples of this phenomenon through percolation, for example. However, when included as a design feature, random mixing usually only takes place in a small region (or multiple small regions) of the network modelled.

In the model outlined in Ch. 5, categorically aged agents (either adults or children) are grouped into households with overall size distribution similar to that of 2016 Statistics Canada census data [96]. Within each household, all agents have effective physical interaction with each other, thereby making each household a clique. As dictated by parameter values, children and adults are enrolled in a sole educational institution, where they interact in assigned classrooms, but also randomly mix in hallways, bathrooms, parking lots and pickup areas, and other such communal areas. In this way, each classroom (and the institution in general) can also be treated as a clique. A simplified example of such a network is shown in Fig. 1.3, where five central mutually connected




Figure 1.3: Representation of the graph structure of an [ABM](#) featuring random mixing within specified cliques. The agents of the model are represented by brown nodes, all of whom interact with every other agent within their enclosing (dashed) polygon. Solid edges between polygons represent the links between cliques by which information is shared.

polygons represent classrooms within which agents (brown nodes) mix randomly. Polygons outside the central cluster represent households of interacting agents, some of whom house school attendees. Further details of the model are outline in [Ch. 5](#), but we can see that random mixing is not necessarily a simplifying assumption, nor is it necessarily an emergent phenomenon caused by assortativity.

1.1.4 Behaviour and graph theory

Contemporary disease spread drives the provision and consumption of information related to vaccination and other public health interventions. If a vaccine-preventable disease is uncommon, more time is available to examine the benefits and perceived costs of the vaccine [403, 268]. Unfortunately, low disease visibility steeply discounts the benefit of vaccination, and fears of vaccine adverse effects go uncontested. This leads to underestimation of the risks of the disease due to its perceived low prevalence or mortality rate, and adoption of an anti-vaccination sentiment based on irrational and ill-informed analyses [248, 405, 474, 413]. Also tied to low disease prevalence is the overestimation of health detriments of vaccination [413, 362, 505], and the occurrence of vaccine scares [254, 43, 212].

These mistaken beliefs about vaccination can also stem from the consumption of information from dubious sources [403]. For example, sensationalisation of vaccine adverse effects can lead to significant overestimation of the danger involved and rampant anti-vaccination sentiment [574, 205]. There are also the prospects of mistaken scientific validity and false consensus [150, 226, 109]. As expected, these ideas persist long after disproof; sometimes, even retraction takes a long time [257, 258, 378, 528, 402].

Interpersonal communication and connectivity are also an important channel of information spread, even more so than media consumption [9, 568]. As information diffuses through a population, it becomes tainted by the personal biases of those spreading it. This is partly due to the psychological theory of reinforcement, which claims a preference for information conforming to views already held in the avoidance of cognitive dissonance [280]. As such, distinct groups based on shared sentiment tend to form in all social interactions [124, 356, 88, 356]; this development of ideological factions is becoming more pronounced as the creation, travel and consumption of media and information becomes easier [198, 176, 48].

Existing within these groups, *echo chambers* are subgroups in which each node has only neighbours with the same attribute (a rigorous definition is given in Sec. 3.2.2). In models of decision/sentiment based on the flow of information and near-neighbour interaction, echo chambers serve as regions of self-reinforcement of sentiment due to the inherent “insulation” of members from exposure to opposite opinion and the ease of information flow [216]. They epitomise the adage “*birds of a feather flock together*” and are fixtures in social networks such as Twitter and Facebook [155].

Salathé et al. confirmed that belief is more important to disease spread than the simple availability of vaccines, with the influence (on the dynamics) exerted by groups of like-minded people becoming more pronounced as vaccine coverage approaches the herd-immunity level³ [478]. A

³Herd immunity refers the disappearance of a disease from a population with sufficiently high vaccination rate; that is, not everyone need be vaccinated. The unvaccinated are insulated from the disease by the rest of the population.

similar study by Fu et al. shows that an adaptive imitation dynamic for vaccination can lead to a catastrophic drop in overall vaccination rates (even below the optimal personal level), as vaccine cost crosses a critical threshold [209].

The models of information and communication presented so far are relatively simple, using adoption or discrete degrees of bias to ‘incorporate’ new information into current beliefs. However, using a sentiment-incorporation rule based on Dempster-Shafer theory⁴ showed that the reporting rates of events in disease and vaccination affected vaccine uptake, while information decay influences the average time taken to vaccinate [567]. Information decay and fidelity is also a major factor in information dissemination and has been demonstrated in social networks (for example, the evolution and spread of memes) on Facebook [7].

In the scope of this project, two network structures are used for the social network: the random graph and the small world network. Random graphs are used in many studies of social interaction[388], and are easily described: the Erdős-Rényi graph (ER, [177]) $\mathcal{G}(N, p)$ is a graph of N nodes constructed by connecting each newly added node to every node already present with probability p . Despite their frequent exploitation, many authors have pointed out shortcomings of the random graph construct that decrease its ‘realism’ and applicability to empirical study [388].

Many studies have however established that empirical networks have highly skewed degree distributions (due to assortative mixing), among other properties such as a high degree of clustering and small average path length. These properties occur in the small world network model proposed by Watts and Strogatz [551]. In this construction, the network $\mathcal{WS}(N, K, \beta)$ with size N , mean degree K and rewiring probability β is made by first building a ring lattice, where each new node is connected to $\frac{K}{2}$ neighbouring nodes on each side. Then, each node is visited and each edge between it and its left neighbours is rewired to some random node with probability β ; this random rewiring drastically reduces the average path length and diameter of the network.

Varying the parameter β allows variance of the network properties between a ring lattice $\beta = 0$ and an approximation of the ER graph $\mathcal{G}\left(N, \frac{K}{N-1}\right)$ when $\beta = 1$. This static model is the second network structure used in this thesis (Ch. 4).

1.2 Critical transition and the “epidemic”

In general, epidemiological literature approaches the study and incidence of epidemics in a great variety of ways; in this thesis, we focus on the phenomena surrounding critical transitions. These

⁴The Dempster–Shafer theory of belief functions is based on Bayesian subjective probability and provides for calculation of the level of belief in a premise (given the belief in a related principle), as well as giving a rule for reliably combining those beliefs so the evidence of belief may be taken from other disjoint sets [70, 576, 348].

[critical transitions](#) occur when some small change in (one of the parameters of) a physical system causes an abrupt change from one state to another as some [critical point](#) is passed [486, 484]. Epidemics can be seen as special cases of critical transitions in the general health of a population [165, 81]. More specifically, [165] casts the recurrence of vaccine-preventable disease as a result of a critical transition in which a change in vaccine uptake bolsters the infection spreading process to some lasting point of sustained transmission.


Critical transitions in epidemiological models have long been related to phase transitions in statistical mechanics [431]. *Phase* refers any state of a system with distinct macroscopic characteristic characteristics in which changes in properties are continuous (and is usually characterised by the value of an order parameter), while *phase transition* occurs when the free energy of the physical system (or one of its derivatives) develops a singularity [431, 573, 247]. For instance, one of the well-known conditions for epidemic spread of the disease is that the [basic reproductive ratio](#) must exceed one $\mathcal{R}_0 > 1$ (the epidemic threshold, seen in Fig. 1.4a) [24]. This is similar to the water phase transitions, in which critical temperatures and pressures must be passed to give the solid, liquid and gas phases (Fig. 1.4b) [501].

Another perspective involves percolation. Many studies have modelled disease as network percolations [285, 540, 387, 369, 386, 481, 363], where an epidemic occurs when a percolation threshold is passed and a giant connected component appears [365, 492]. That is, when the size of the largest cluster of susceptible agents “becomes comparable” to the size of the population [369], or more specifically, the size of the component varies with the size of the model (population) [247, 286, 352]. Another more concrete yet equivalent statement of an epidemic is given by [386] as


“Epidemics are defined as outbreaks that affect a non-zero fraction of the population in the limit of large system size.”

This speaks to the threshold behaviour treated in Def. 3 in that this is when the proportion p of susceptible agents in the model passes some calculable percolation threshold p_c [369, 281] (in many sources, it is outlined as a condition *conducive to* [epidemic](#) spread, with any remaining interpretation left to the reader). This demonstrates the physical concepts of scale invariance and [correlation length](#) divergence that accompany a [second-order](#) phase transition in physical systems [432]. [486] indicates that such a continuous phase change may be due to the heterogeneity and relative low connectivity.

Despite the comfort offered by this definition, it relies on the (manipulation of the) property of network size, an oft-seen theme in this research. Many following research papers (especially those featuring computational models) obtaining results from [ABMs](#) speak of “epidemic spread” while giving no clear description of what the term ‘epidemic’ means within context despite arguable necessity, a phenomenon well noted in [414]. Based on the descriptions of graphs presented in these papers and references in their discussion sections, the author can only conclude that the



(a) Final size of an epidemic with respect to the basic reproduction \mathcal{R}_0 .



(b) Phase change diagram of water (shown with respect to pressure and temperature).

Figure 1.4: Comparing the epidemic threshold and the phase transitions of water.

(a) “Epidemic phase transition. Final size of an epidemic as a function of its basic reproductive ratio \mathcal{R}_0 , for a susceptible-infected-recovered (SIR) model with a homogeneous network structure, with a number of connections (k) of 4 for each individual. Transmission rate β varies between 0 and 3 with recovery rate $\gamma = 1$, resulting in \mathcal{R}_0 ranging between 0 and 3. The line depicts the analytical results whereas the red dots show the results from stochastic simulations with a population size of 104. The epidemic does not occur for $\mathcal{R}_0 < 1$, whereas the final size increases as a function of \mathcal{R}_0 for values higher than 1. The analytical results and the simulations are in good agreement.” [179]

(b) author of the original work: Cmglee / CC BY-SA (<https://creativecommons.org/licenses/by-sa/3.0>)

word ‘epidemic’ in this sense (independent of network size) is a loose idea of ‘disease explosion’ in a large group of susceptible agents, as opposed to the ‘slow burn’ seen in populations where most of the population is in some way immune to the infection. These ideas of speed and spreading rate partially inspire the description given in Def. 1.2.

Figure 1.5 shows a graphic representation of an epidemic. A blue foreground bar chart shows the trend in the proportion of infected individuals in the population over time, with a grey-shaded region representing a family of curves generated by some EBM of the infection process. The pink-shaded regions show the areas of the trend that would be termed as epidemics in the loose sense: the lead up to a peak region in the infection trend of the model. In comparison, the homogeneous-mixing assumption made in EBMs is (in this simple sense) equivalent to a spreading process unfolding on a *complete network*, a contact network in which each agent has effective contact with every other agent in the population. This is precisely the methodology used in Ch. 5, in which we treat the subnetwork of school attendees as being complete, with each




Figure 1.5: A contrived representation of a family of curves giving the proportion of infected agents (grey shaded area), with the true trend given by the blue bar chart. Pink-shaded regions represent what would be termed an ‘epidemic’; that is, a region of the trend indicating fast growth of infection in the population being modelled.

agent having effective contact with everyone else during common area interactions.

Here, we can define an *epidemic* as a large infection spreading process distinct in size from any established endemic infection rate [290, 200, 493, 287], or as a bifurcation [486, 91, 334, 108, 30, 13] ending in an endemic regime [414], both confined to a well-defined region being investigated. For our purposes, the characteristics of an epidemic will be as follows:

1. the transmission and/or acquisition of a common (strain of the) disease occurs through effective agent-to-agent contacts or some other common vector [200, 287],
2. the time taken for a sudden wave of infection to weaken is small relative to the length of time needed to establish either endemicity or absence of the same disease in the region specified [414],
3. there exist thresholds in the dynamic for either network and compartment size [41, 384, 282, 222, 526], household reproduction number [437], transmissibility [369], average neighbourhood size [455] or behavioural changes [565, 360, 477],
4. the spread of the disease is limited only by the number of susceptible agents on the network; given that the transmissibility of the disease exceeds some threshold value, the infection dies only due to the (near-complete) depletion of the pool of susceptible agents (rather than lack of transmission) [290, 200, 55],
5. the peak number of agents infected during the spread is on the order of the size of the population under study [437],

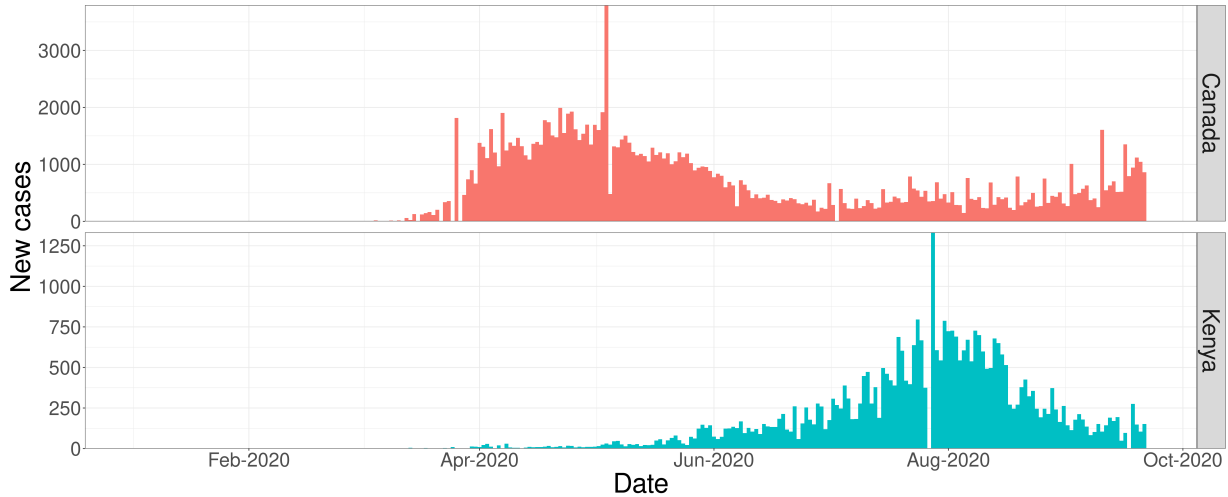


Figure 1.6: Plots of the number of new [COVID-19](#) cases per day for Canada and Kenya, from January to September 2020.

6. the spreading rate of the infection is superlinear (due to the positive feedback loops formed by interactions between infected and susceptible agents) [113, 290, 333]. Alternately, the velocity of the epidemic can be measured by established methods given in existing literature [535, 120, 121].

A more rigorous treatment of epidemics (and their definition) is given by various sources [290, 414], while Fig. 1.6 shows a typical epi curve: the number of new [COVID-19](#) infections per day in Canada and Kenya between January and September 2020 spread through community infection [417].

The approach of these critical transitions can be typically accompanied by many characteristic phenomena such as flickering [484, 549], critical slowing down [484, 486], critical speeding up [531, 457, 76], skewness [484], autocorrelation [139, 138], variance [138, 301] and others [486]. In this thesis, we focus on [EWS](#) resulting from [critical slowing down](#), which occurs when a system takes progressively longer to recover from a small perturbation (due to weakened stability of the current system equilibrium) as some critical point is approached [80]. Specifically exploiting the network structure of the models outlined in Sec. 1.3, we will focus on spatial indicators.

1.3 A coupled behaviour-infection model

We use an [ABM](#) in Chs. 2, 3 and 4 to model heterogeneous mixing among individuals with different opinions and states of health, as well as to impose an [ER](#) network structure on each of

the disease and communication networks; not only does this *bottom-up* modelling approach allow the capture of elusive analytically intractable dynamics not captured by differential equation models [466, 430, 193, 73, 157], but also allows us to view the individual natural history of each agent [215, 472, 73, 107] and track the movement of agents and interactions in dynamic models [157, 28, 433, 169].

We model a multilayer network where each layer is given an identical undirected ER random graph with size N and mean node degree $\langle d_n \rangle$. Each agent n can be described by a pair of states; for instance, each agent is assigned the joint state (V_s, V_p) at the start of the simulation with probability α (they are a pro-vaccine vaccinated agent), else they are initialised with joint state (N, S) (an anti-vaccine susceptible agent) with probability $1 - \alpha$.

The social process follows an NV_s dynamic (Fig. 1.7b), representing pro- (V_s) and anti-vaccine (N) opinion for each agent n . ξ represents the probability of any agent switching opinion randomly in each week and $\mathbb{P}_n(N \rightarrow V_s)$ represents the probability of switching from anti-vaccine opinion to pro-vaccine opinion ($N \rightarrow V_s$) upon interaction with a disagreeing neighbour. We introduce an imitation dynamic by having each agent n compare its opinion with a single randomly chosen social contact (a neighbouring agent on the social layer) each week; n then changes its vaccination opinion only if there is disagreement (the agent and the neighbour have different vaccine opinions). This change of opinion depends on the perceived risk of vaccine adverse effects κ (“*perceived vaccine risk*”) and I_n (the number of infected physical neighbours of n) according to the rules


$$\begin{aligned}\mathbb{P}_n(N \rightarrow V_s) &= \frac{1}{1 + \exp\left(-U_n^{N \rightarrow V_s}\right)}, \\ \mathbb{P}_n(V_s \rightarrow N) &= \frac{1}{1 + \exp\left(-U_n^{V_s \rightarrow N}\right)},\end{aligned}\tag{1.1}$$

where the indices $U_n^{N \rightarrow V_s}$ and $U_n^{V_s \rightarrow N}$ in (1.1) are utility functions defined as


$$\begin{aligned}U_n^{N \rightarrow V_s} &= -\sigma \left(\frac{d_n^N - d_n^{V_s}}{d_n} \right) - (\kappa - I_n), \\ U_n^{V_s \rightarrow N} &= -\sigma \left(\frac{d_n^{V_s} - d_n^N}{d_n} \right) + (\kappa - I_n).\end{aligned}\tag{1.2}$$

Here, σ represents the strength of an injunctive social norm to maintain the current opinion (referred to as “*social norm*”), while d_n^* represents the number of neighbours of n with opinion $*$ and d_n represents the total number of social contacts of n .

The epidemiological dynamics (which we term ‘infection dynamics’) follow an $SIRV_p$ process (Fig. 1.7a), in which an agent n can progress through each of four disease compartments: S (susceptible), I (symptomatically infectious/infected), R (recovered) and V_p (vaccinated). Every



(a) Schematic of the infection dynamics of the model. Effective contacts occur between susceptible S and infected I agents with probability p per time step (1 week). Upon deciding to vaccinate (with probability $\mathbb{P}_n(N \rightarrow V_s)$), a susceptible agent n becomes physically vaccinated ($S \rightarrow V_p$). Infection lasts $\ell = 2$ weeks after which agents recover ($I \rightarrow R$). Upon death (with probability μ per week), an agent is “reborned” with either vaccinated (probability $\alpha \cdot \mu$) or susceptible (probability $(1 - \alpha) \cdot \mu$) status.



(b) Representation of the opinion dynamics of the model. Per time step, each agent switches between pro- (V_s) and anti-vaccine (N) opinion with probabilities $\mathbb{P}_n(N \rightarrow V_s)$ and $\mathbb{P}_n(V_s \rightarrow N)$ respectively upon interaction with a dissenting neighbour. α gives the probability of being birthed with pro-vaccine opinion V_s .

Figure 1.7: Diagrams showing the physical (a) and social (b) dynamics of model V1.

week (i.e. in each time step), each susceptible agent n interacts with all its physical neighbours; each effective interaction carries the probability p of infection ($S \rightarrow I$), so that each susceptible agent faces the total probability $1 - (1 - p)^{I_n}$ of infection in a single week. The duration of the illness is ℓ weeks (with no impact on mortality), after which n gains lifelong natural immunity

$(I \rightarrow R)$.

Alternately, if a susceptible agent n adopts a pro-vaccine opinion, they are immediately vaccinated ($S \rightarrow V_p$) and gain lifelong artificial immunity. We also assume that only susceptible agents are vaccinated. Thus, individual agents may change their opinion about vaccination multiple times in their life ($N \rightarrow V_s \rightarrow N$), but once they are vaccinated they may not become unvaccinated ($S \rightarrow V_p$). This in turn creates an asymmetry between disease dynamics and social dynamics that will have implications for the model predictions.

Each agent n has probability μ of dying each week, upon which they are replaced by a new pro-vaccine vaccinated agent (V_s, V_p) with probability α , or an anti-vaccine susceptible agent (N, S) with probability $1 - \alpha$, keeping the same physical and social contacts as the agent they replaced (so that the network is static). Case importation is accounted for by infecting a randomly selected proportion ι of susceptible agents at the start of each week, and noise is introduced to the model by changing the vaccine opinions of a randomly selected proportion ξ of the entire population weekly.

At the start of the simulation, some susceptible agent is randomly selected as an index patient and infected; subsequent disease spread is governed solely by environment and agent-agent interaction. Figure 1.8 gives a flowchart demonstrating the flow of agents through the different phases of the model occurring every time step (for most of the processes). Processes in yellow boxes occur only once throughout the *realisation* while processes in red boxes denote loops, where the instruction is repeated for all agents in the network. Blue diamonds represent true/false decisions, and grey boxes represent choices. The steps are:

1. The realisation is initialised with a ratio α of the agents assigned social state V_s and physical state V_p (i.e., vaccinated pro-vaccine agents), and the remainder assigned the physical state S and social state N (i.e., susceptible anti-vaccine agents). Both layers of the network are given identical random network structure, so that each agent's social contacts are also their physical contacts and *vice versa*.
2. A single susceptible agent is chosen as the index patient of the disease, and infected ($S \rightarrow I$).
3. Every susceptible agent in the network interacts with *all* of its physical neighbours, and every such contact carries probability p of infection spread. Since the agent n has I_n many physical contacts, n 's total probability of infection in a single time step will be $1 - (1 - p)^{I_n}$.
4. Every time step, each agent n in the network compares opinion with some random social neighbour a , with opinion change only if they disagree (different social states). If they do disagree, then agent n will adopt a 's vaccine opinion with probability P_n , and is immediately vaccinated if susceptible *and* changing to a pro-vaccine opinion. The next agent in the network is selected, and the loop is repeated until there are no more agents to be considered.

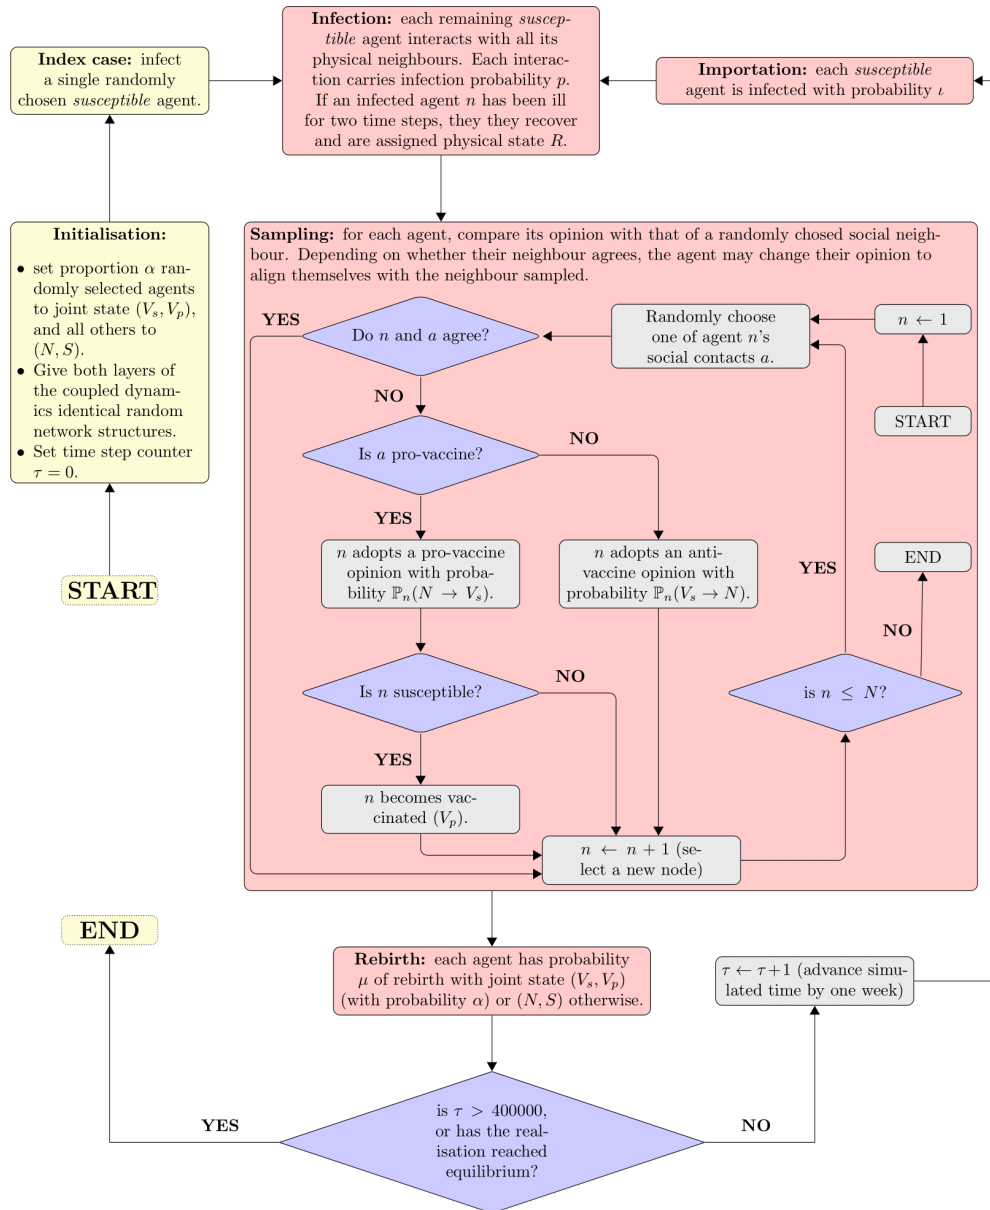


Figure 1.8: Implementation of the model dynamics used for each stochastic realisation. Red blocks represent loops run for each node in the network per time step. Yellow blocks are executed only once per realisation. Blue diamonds represent binary decisions, and grey boxes represent simple instructions.

5. In each time step, every agent faces probability μ of death, at which they are instantly replaced by either a vaccinated pro-vaccine agent (joint state (V_s, V_p)) with probability α , or a susceptible anti-vaccine agent (N, S) .
6. Then if the **realisation** has converged, it ends and the state of the system is output for each time step run; else, it progresses to the next time step.
7. The first process of every subsequent time step is **case importation**, where proportion ι of susceptible agents is infected before the infection process is repeated in the simulation.

Table 1.2 lists the variables and notation used in Chs. 2, 3, and 4.

State	Description
S	susceptible to infection
I	infected and infectious
R	recovered from illness, immune until death
V_p	vaccinated, immune until death
N	has an anti-vaccine stance
V_s	has a pro-vaccine opinion
Notation	Description
I_n	number of infected neighbours
$\mathbb{P}_n(\lambda \rightarrow \rho)$	probability of switching sentiment from λ to ρ , given in (1.1)
d_n	number of neighbours of n
d_n^λ	number of neighbours of n with sentiment λ
$U_n^{\lambda \rightarrow \rho}$	utility of switching from sentiment λ to ρ , given in (1.2)
$[\lambda]$	proportion of agents with vaccine opinion λ at some time step τ
$\langle \lambda \rangle$	mean value of $[\lambda]$ over all realisations
$[\lambda, \rho]$	number of edges between agents with vaccine opinions λ and ρ
$\langle \lambda, \rho \rangle$	mean value of $[\lambda, \rho]$ over all realisations

Table 1.2: Table of all variables appearing in Chs. 2, 3, and 4.

Since all the models used in this thesis feature undirected networks; $[\lambda, \rho] = [\rho, \lambda]$ for sentiments λ and ρ . Table 1.3 gives a list of the parameters used in all the models of Ch. 2. All networks were created using the Stanford Network Analysis Project (SNAP) library (for C++) [89].

Parameter	Description	Unit
N	number of agents	constant
ι	case importation rate	wk ⁻¹
p	disease infectivity	constant
ξ_*	probability of randomly changing sentiment	wk ⁻¹
σ	strength of the social norm	constant
κ	perceived risk of vaccine adverse events	constant
\mathcal{T}	length of each realisation of parameter values	constant
α	initial ratio of pro-vaccine vaccinated agents	constant
ℓ	duration of the infection	constant
μ	birth/death rate	wk ⁻¹

Table 1.3: Table of all parameters appearing in Chs. 2, 3 and 4.

1.4 Equilibrium and stopping criteria

We specify that any realisation of parameter values has reached a (*computational*) *equilibrium* when the output variables of the simulation over the last 500 time steps all have a standard deviation of 0.05% of their respective maximum values.

In the models used in this thesis, the output variables (called *model variables*) are each of $\langle V_s \rangle$ (number of pro-vaccine agents), $\langle N \rangle$ (number of anti-vaccine agents), $\langle S \rangle$ (susceptible agents), $\langle I \rangle$ (infected agents), $\langle R \rangle$ (recovered agents), $\langle V_p \rangle$ (vaccinated agents), $\langle N, N \rangle$ (links between anti-vaccine agents), $\langle N, V_s \rangle$ (dissimilar joins) and $\langle V_s, V_s \rangle$ (links between pro-vaccine agents). Our model intrinsically satisfies conditions (1), (2), (4), (5) and (6) of the definition in Sec. 1.2.

1.5 Change point testing

Some previous studies proffer EWS of disease spread without attempting to rigorously verify two *snuck premises*: 1) the EWS give (computationally recognisable) warning signals, and 2) these signal points precede some well-defined turning point in the model dynamics; instead, visual heuristic analysis⁵ is used. Besides our calculations of estimates of the transition values in the model dynamics, we employ *change point detection tests* in an attempt to (programmatically)

⁵such rudimentary checks are usually referred to as ‘eyeballing’, but the author would instead like to propose the use of the somewhat more sciencey-sounding technical phrase *visual heuristic analysis* (VHA).

identify and establish times at which warnings are given by the EWS, and a method of comparison of the length of forewarning that we can reasonably expect from these EWS.

The [change point](#) detection tests used in this study are the standard normal homogeneity test (SNHT), Lanzante test, Pettitt test and the Buishand range test; stemming from climatological studies, all test the null hypothesis of homogeneity against the alternative hypothesis of change at one or more identifiable points. Importantly, these four curated tests also allow for the estimation of a location of the change point in the series, unlike other tests such as the von Neumann’s ratio test.

The Lanzante, Pettitt and Buishand range tests are all non-parametric (they make no assumption of the underlying distribution of the data) and they test the null hypothesis H_0 of no shift of the central tendency of the time series tested. Though Pettitt’s test is held as the most commonly used change point detection test [270, 319, 441, 90] and it was used in Chs. 2 and 3 (also used for comparison of results in Ch. 4, though not shown), we chose to primarily feature other tests due to inconsistencies and difficulties with lead distance results (and even the presence of κ -series change points) obtained from the Pettitt tests. Exploration as to exactly why these difficulties arose was considered to be beyond the scope of this project; nonetheless, the optimal application of these change point tests to the generated κ -series remains of interest to the author.

The SNHT compares the standardised ratios of the observations of the first n values of the series with the mean of all other observations, and tests the null hypothesis of a $\mathcal{N}(0, 1)$ distribution of these ratios against the alternate hypothesis of a shift in the location of the distribution to some $\mathcal{N}(\mu, 1)$ [14, 426].

1.6 Objectives

The objective of this thesis is to identify a class of dependable early warning signals of vaccine crises and epidemic events in coupled disease–behaviour models of paediatric infectious diseases. In working towards this, there are two sets of research questions.

First, can we find a class of statistical tools to predict critical transitions in coupled behaviour-infection models? If so, can we expect these tools to be of general use? If not, can we formulate criteria for choosing suitable early warning signals based on recognisable properties of the system in question? Additionally, how do the various EWS compare with each other? Here, we formulate an idealised model of childhood disease, and establish the response of some proposed tools.

Secondly, having found (and/or formulated) these warning signals, do the EWS fail in any reasonably foreseeable contexts? If so, can these weaknesses in the proposed tools be overcome? We will use a combination of parameter searches, changing network structures, and gradually varied network properties to probe for “breaking points” of the remaining tools. This will aid us in understanding how these early warning signals should be properly applied.

1.7 Outline

In line with our objectives, the thesis will proceed as follows. Our first study (Ch. 2) will introduce the first set of **early warning signal (EWS)** we will explore. Epidemic events and convergence will be defined, and used throughout the remainder of the thesis. Simulations run on variations of the **ABM** outlined in Sec. 1.3 (named **V1**, **V2** and **V3**) will show that mutual information, join counts statistics, Moran's I and Geary's C can act as EWS of the system. The method of application of the change point tests used will be tested, with the resulting method applied to the the model data. Lead distances of each EWS will allow comparison among the EWS, and between the EWS and raw model outputs. We will demonstrate that Moran's I and Geary's C are simple linear combinations of join counts.

Our second study (Ch. 3) will add community and neighbourhood structure, the probability of having an infected neighbour, opinion network clustering coefficient and modularity score to the list of EWS. The concepts of community and echo chamber will be clearly defined. Simulations run on model **V1** (from Ch. 2) will show that these newly introduced tools are also stable EWS of the transitions of the model dynamics. The new EWS will be compared via change point tests.

The third study in this thesis (Ch. 4) will add new EWS (network diameter and triad census) to the group. A new **ABM** (named **V4**) will be introduced that features a changed network structure, the inclusion of vaccine hesitance and data gathering through sampling. To accommodate the changes brought by **V4**, key EWS measurements will be redefined. Simulation data from **V4** will be used to compare the results to the previous studies (Chs. 2, 3), and the effects of the changes in the model on the performance of the EWS will be explored.

Our fourth study (Ch. 5) models the transmission of **COVID-19** in childcare centres and primary schools in Ontario. Recent literature on the transmission of **COVID-19** will be discussed at length. Simulations run on a model population of households and a single school will investigate the number of produced infections in a combination of reopening scenarios varying classroom size, classroom arrangement and cohorting. The location with greatest infections and the number of student-days lost due to classroom closure will be investigated, with sensitivity analysis given in App. 5.

Finally, Ch. 6 will collate the results of the studies and discuss the successes and limitations of the the research. Findings will be contextualised using existing literature in the field, and broader implications discussed.

Chapter 2

Spatial aggregation can indicate regime shifts

Abstract

The resurgence of infectious diseases due to vaccine refusal has highlighted the role of interactions between infection dynamics and the spread of vaccine opinion on social networks. Shifts between infection elimination and outbreak regimes often occur through tipping points. It is known that tipping points can be predicted by early warning signals (EWS) based on characteristic dynamics near the critical transition, but the study of EWS in coupled behaviour-infection models has received little attention. Here, we test several EWS indicators measuring spatial coherence and autocorrelation for their ability to predict a critical transition corresponding to disease outbreaks and vaccine refusal in a multiplex network model. The model couples paediatric infectious disease spread through a contact network to binary opinion dynamics of vaccine opinion on a social network. Through change point detection, we find that mutual information and join count indicators provided the best EWS. We also show the paediatric infectious disease natural history generates a discrepancy between population-level vaccine opinions and vaccine immunity status, such that transitions in the social network may occur before epidemiological transitions. These results suggest that monitoring social media for EWS of paediatric infectious disease outbreaks using these spatial indicators could be successful.

Material in this chapter is based on the publication:
Phillips, B., Anand, M. & Bauch, C.T. Spatial early warning signals of social and epidemiological tipping points in a coupled behaviour-disease network. *Sci Rep* 10, 7611 (2020). <https://doi.org/10.1038/s41598-020-63849-0>.

2.1 Introduction

Resurgences of vaccine-preventable diseases can severely stress public health systems [561], interrupt tourism and public services [31] and decrease GDP [34] through the huge costs of large-scale interventions [231, 350, 366]. Some of these events have been driven by falling vaccination rates due to vaccine refusal [443], for example the persistence of endemic polio in countries such as Nigeria [105, 469, 156]. Intentional undervaccination may be attributed to complacency [245, 408] and the spread of anti-vaccine sentiment facilitated by media coverage and its sensationalisation of true adverse vaccine effects [574], misstatement of the cause of illnesses [205], the spread of rumours and false information [533, 506] and the effect of social norms [508]. False reporting of adverse effects can also be reinforced simply by the expectation of such side effects [504].

These phenomena illustrate how the social diffusion of information is heavily responsible for the trajectory of infection spread through its ability to alter individual behaviour. Much work has modelled opinion dynamics for different applications [502, 269, 487], where choice is modelled by stochastic differential equations [147], voter models [332, 143], Markov chains [152], near-neighbour averaging [207], majority opinion models [210, 298], impact models [401] and flocking models [330]. The combination of these frameworks with network structure has revealed much about the occurrence of opinion cascades [271, 284, 229] and forecasting opinions [147, 289]. Models coupling behavioural dynamics and spreading processes have yielded useful results, such as the ability to predict the risk of illness of an individual [132], reasonable predictions of influenza transmission patterns [83, 436], price dynamics [468], infection outbreaks [274] and predictions of climate change [52].

Opinion propagation can be intuitively represented by diffusion of information through social networks [125, 565, 360]. Similarly, infection spread is often conceptualised as spreading through a physical contact network [281]. An increasing number of models explores the dynamics of *n-layer multiplex networks*, where each layer represents a different aspect of the dynamics of a single coupled system. For instance, a growing body of work studies coupled behaviour-infection dynamics on 2-layer multiplex networks, where social dynamics propagate through a social layer and infection dynamics propagate through an infection layer [44, 17, 566, 584, 274]. In these cases, the theory of *phase transitions* in spatially structured systems is important; for instance, epidemic regimes have previously been modelled as the outcome of phase transitions in physical systems [272, 570, 288, 262, 364, 452, 249, 165, 173, 313]. Generally, a *phase transition* occurs when a physical system moves from one state to an alternate state [385]; the point at which this transition occurs is called a *critical point*. *First-order* transitions occur when any macroscopic variable varies discontinuously (for example, a jump between vaccinated and non-vaccinated regimes in an infection model), while *second-order* transitions occur when a macroscopic variable varies continuously [195, 170]. These second-order transitions are also called *critical transitions* [404, 234] and have seen much interest in the epidemiology literature [165, 484, 485, 486, 415, 416, 173, 164, 302]. Systems approaching these critical transitions sometimes display characteristic spatial or temporal

behaviours called *early warning signals* (EWS) that can help to predict coming epidemic infection outbreaks and other events [301]. EWS can be defined as statistically significant, recognisable, characteristic behaviours known to precede critical transition in dynamical systems [81, 484, 485]; some are manifestations of *critical slowing down*, where the relaxation time (the time taken by the system to return to an equilibrium state after a small perturbation) increases as a *critical point* is approached [484, 415]. Critical slowing down can be shown to precede both first- and second-order transitions (as well as other types) [485, 562, 517, 144, 484] and is accompanied by the divergence of *correlation length* in a physical system [144].

A variety of statistics have been used to study early warning signals in spatially extended systems. For instance, temporal correlation [138, 312, 415, 471, 164, 484, 49, 162, 46] and spatial correlation [139, 163, 274, 484, 420, 283] have been found to precede transitions in spreading processes, including infection processes. Other metrics have been applied to spin systems: each site in a lattice may be in one of two possible states and the probability that a site is in a given state depends partly on the state of neighbouring sites. The spin model has also been applied to opinion dynamics; a simple voter model with binary opinion dynamics is analogous to a physical spin system, where particles represent agents and positive and negative spins represent two types of opinions [523]. Consensus formation can then be seen as a second-order phase transition to an ordered (magnetic) state, where all spins are aligned; in this regime, knowledge of the opinion of a single agent gives the opinion of all other agents in the system [377]. Since the transition in finite networks is smooth [67, 68], the typical distance over which the opinions of two agents agree increases smoothly; this is analogous to a smooth increase in the correlation length of a physical system [579].

Continuing the metaphor of spin systems, the disordered system takes the form of a *spin glass* above the critical temperature, where there is no particular alignment of the spins of the particles. In a spatial opinion model, this describes a state where opinions between neighbours are generally uncorrelated [545]. On a static network, this state should produce a larger number of edges between dissimilar neighbours as compared to that of consensus regimes. This leads to a prediction about the *join count* statistics of the network, where the numbers of edges between like neighbours are compared to the number between dislike neighbours as a test of geographical distribution [371, 196]. This is arguably the most natural (and well-defined) measure for graphs presenting binary data [196]. The statistic saw early use in the 1940s-50s [300, 220, 371] and is currently used for spatial analysis [196], although it has only recently appeared in the early warning signals literature (to the authors' best knowledge).

The necessity of infection surveillance and early warning signals for outbreaks has been discussed in multiple contexts, from epidemic alleviation to bioterrorism [409, 199, 119, 422, 548, 571, 379, 510, 381]. Approaches include tracking migration patterns of disease vectors [194], animal deaths [374], patterns in time series of reported cases [541], monitoring online news sources [206, 85] and searches [456, 182, 264], “bet-hedging” [130] and analysing large data sets derived from social networks [253, 132, 500, 12, 480]. Some tools proposed so far were found to be impractical [253],

while others failed [325] and some have succeeded [57]. Some studies have estimated the benefits of early (and targeted) intervention to be large [184, 557, 353, 341, 509], though others caution against dependence on such predictive approaches [256]. Such potential mitigation of unnecessary expense motivates our interest in finding dependable tools to be used as early warning signals which remain easily computable on large high-resolution data sets. Furthermore, the study of early warning signals in coupled behaviour-infection systems in general (and in two-layer behaviour-infection multiplex networks in particular) has received relatively little attention [425, 274], suggesting a significant gap in the literature.

Our objective is to evaluate and compare the relative merits of the mutual information, Moran’s I, Geary’s C and join count statistics as EWS of the occurrence of epidemics and changes in aggregate opinion on a coupled behaviour-infection network model. We use three differently parametrised models (V1, V2 and V3) coupling a binary vaccination opinion dynamic to an SIRV epidemic process. The resulting trends in the EWS for model V2 will be explored in Secs. 2.3 and 2.4, with V1 and V3 presented in App. A.2.

The outline of this study is as follows: Sec. 2.2 will present the EWS and their derivations and give the details of the model used. Section 2.3 will analyse the trends in the warning signals and Sec. 2.4 will present a review of the study and any shortcomings of our approach, with further results pertinent to the study presented in App. A.

2.2 Methods

We assume an acute, self-limiting infection that confers lifelong natural immunity upon recovery and for which a vaccine is readily available. Similar premises have been used to represent the natural history of many paediatric infectious diseases such as measles [42, 43]. In particular, we assume an $SIRV_p$ natural history consisting of four mutually exclusive epidemiological states. Agents are initially susceptible to infection (S). Upon infection, the agent enters the infected state ($S \rightarrow I$), which we treat as a combination of the latent, ill and infectious periods [21]. Upon clearing the infection, agents enter the recovered state with lifelong immunity ($I \rightarrow R$); additionally, susceptible agents may be vaccinated and so enter the vaccinated state ($S \rightarrow V_p$) [543].

We also include injunctive social norms (i.e. peer pressure) as well as a risk of vaccination that captures both economic costs and the fear of perceived adverse vaccine effects [413]. As in some models [274], we include a noise parameter ξ to account for noise [69] with the simplifying assumption of perfect vaccination [516] (reversion from the recovered state to the susceptible only through agent death). During simulation, each time step represents a single week. Models V1, V2 and V3 all use the multiplex model dynamics laid out in Sec. 1.3.

2.2.1 Early warning signals

Mutual information (\mathcal{M}) is defined as

$$\mathcal{M}(\mathcal{X}, \mathcal{Y}) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \mathbb{P}(x, y) \cdot \log_2 \left(\frac{\mathbb{P}(x, y)}{\mathbb{P}(x) \cdot \mathbb{P}(y)} \right), \quad (2.1)$$

where x and y are discrete random variables; x takes value on the set $\mathcal{X} = \{x_1, x_2, x_3, \dots\}$ and y on set $\mathcal{Y} = \{y_1, y_2, y_3, \dots\}$, with \mathbb{P} a joint probability mass function of X and Y [237]. Mutual information is an entropy-based quantification of the “shared information” of two random variables quantifying how knowledge of one decreases the uncertainty of the other and vice versa [74]. Mutual information peaks at the critical temperature of spin systems during second-order transitions and has been widely used in detecting phase transitions [511, 25]; an advantage of this statistic is its ability to quantify non-linear dependence unlike Moran’s I and covariance, which only account for linear dependence.

Join counts quantify the degree of clustering by giving the number of adjacencies between agents of different types. We divide the population into two attributive classes, with V_s the compartment of pro-vaccine agents and N the compartment of anti-vaccine agents. Let $[\Psi, \Omega]$ be the number of social interactions between agents with vaccine opinions Ψ and Ω ; then $[N, N]$ represents the number of nearest-neighbour interactions between anti-vaccine agents, $[V_s, V_s]$ the number of interactions between pro-vaccine agents and $[N, V_s]$ the number of interactions between pro- and anti-vaccine agents. These can be written as

$$\begin{aligned} [N, V_s] &= \frac{1}{2} \sum_{j,k} \omega_{jk} (x_j - x_k)^2, \\ [N, N] &= \frac{1}{2} \sum_{j,k} \omega_{jk} (1 - x_j)(1 - x_k), \\ [V_s, V_s] &= \frac{1}{2} \sum_{j,k} \omega_{jk} x_j x_k, \end{aligned} \quad (2.2)$$

where $\omega_{j,k} = 1$ if agents j and k are social neighbours and $\omega_{j,k} = 0$ otherwise (ω is the adjacency matrix of the social network); x_n represents the *opinion score* of agent k , defined as

$$x_k = \begin{cases} 1, & k \in V_s \quad (x_n \text{ has a pro-vaccine opinion}) \\ 0, & \text{else} \end{cases}. \quad (2.3)$$

In an opinion model, clustering is manifested by agents consistently having a higher number of like-minded neighbours than expected based on the global prevalence of the opinion; join counts are then used to test the null hypothesis of positive correlation [201]. Join counts are

used in many fields as a categorical test of spatial autocorrelation, including ecology [453] and geographical information systems [22]. In all parameter realisations here, the number of joins are counted naïvely rather than calculated. Joins between like-minded agents (e.g. $[N, N]$ and $[V_s, V_s]$ joins) will be called *similar joins* and edges between disagreeing neighbours (e.g. $[N, V_s]$) will be called *dissimilar joins*.

The *Moran's I coefficient* \mathcal{I} quantifies spatial correlation and is defined

$$\mathcal{I} = \frac{D}{W} \cdot \frac{\sum_{j,k} \omega_{jk} (x_j - \bar{x})(x_k - \bar{x})}{\sum_j (x_j - \bar{x})^2}, \quad (2.4)$$

where the number of edges W is given as $W = \sum_{j,k} \omega_{jk}$, D is the size of the network and $\bar{x} = \frac{1}{D} \sum_j x_j$ represents the mean opinion score of the population [196]. Used as a global statistic, Moran's I gives the degree of correlation between the values of neighbouring patches (agents and their social neighbours); here, the numerical value of the vaccine opinion is the same as described in (2.3). Algebraic manipulation of (2.4) using (2.3) gives

$$\mathcal{I} = \frac{D}{W} \cdot \frac{2 \cdot [V_s, V_s] - 2\bar{x} \cdot (2 [V_s, V_s] + [N, V_s]) + W \cdot \bar{x}^2}{(1 - 2\bar{x}) \cdot [V_s] + D \cdot \bar{x}^2}, \quad (2.5)$$

(full derivation given in App. A.1); we can then consider Moran's I as a measure derived from the linear combination of join counts. Positive values signify spatial correlation, with negative values signifying anti-correlation.

The *Geary's C coefficient* \mathcal{C} is yet another measure of spatial correlation based on the cross-product (like Moran's I) [196], but unlike Moran's I it accounts for the difference in opinion between two neighbours [221]. It is given as

$$\mathcal{C} = \frac{D - 1}{W} \frac{\sum_{j,k} \omega_{jk} (x_j - x_k)^2}{\sum_j (x_j - \bar{x})^2}. \quad (2.6)$$

Lower values show spatial correlation and large values represent anticorrelation. Like Moran's I (2.5), Geary's C can also be expressed as a linear combination of join counts

$$\mathcal{C} = \frac{D - 1}{W} \frac{2 \cdot [N, V]}{(1 - 2\bar{x}) \cdot [V] + D \cdot \bar{x}^2}; \quad (2.7)$$

this expression is also derived in App. A.1.

2.2.2 Parametrisation

The birth/death rate in the model was set at $\mu = 2.4 \times 10^{-4}$, giving each agent a mean life expectancy of 80 years. The network size $D = 40000$ was chosen to represent a small town

where each agent n has effective physical contact with $\langle d_n \rangle = 30$ neighbours per week, where an **effective contact** is defined as any interaction between agents that allows for infection and/or communication of opinion. The case importation ratio $\iota = 2.5 \times 10^{-5}$ was added to provide periodic impulses of infection as a test of resilience in endemic disease regimes. Here, an ensemble of 100 simulations using parameters $\kappa = 0$, $\sigma = 0$ and $\alpha = 0.05$ returned the values $\langle S \rangle, \langle R \rangle < 0.05$ at *equilibrium* (defined in Sec. 1.2), where $\langle \Psi \rangle$ represents the mean number of agents with (social or epidemiological) state Ψ , averaged over all realisations of that combination of parameter values.

The infectivity $p = 0.2$ was chosen to reflect the **basic reproductive ratio** of a measles infection commonly estimated from empirical data [238]; effective physical contacts occur in the simulation once per week during the period of infection. The probability of randomly switching opinion $\xi_1 = 1 \times 10^{-4}$ was included as a source of noise. We found that the parameter ranges for perceived vaccine risk $\kappa \in [-1, 1]$ and social norm $\sigma \in [0, 3]$ were sufficiently broad to capture transitions in both social (Fig. 2.1a) and infection dynamics (Fig. 2.1b), as well as the corresponding trends in the mutual information (Fig. 2.1c) and dissimilar join count (Fig. 2.1d) statistics.

The contours in each panel of Fig. 2.1 show the obvious correspondence between transitions in the social (Fig. 2.1a) and infection (Fig. 2.1b) dynamics of the model, and substantial changes in $\langle \mathcal{M} \rangle$ and $\langle N, N \rangle$; here, the dissimilar join count $\langle N, V_s \rangle$ (Fig. 2.1d) and mutual information \mathcal{M} (Fig. 2.1c) increase while the perceived vaccine risk κ and social norm σ parameters increase towards their respective (pre-transition) threshold values. These trends are generally asymmetric about both transitions; this can be seen in Fig. 2.3, where post-transition trends do not exhibit similarly detectable warnings (if any). This parametrisation applies to model V2. The corresponding parametrisations and contour plots of models V1 and V3, as well as their post-transition trends are presented in App. A.2.

2.3 Results

Due to the low initial vaccine coverage $\alpha = 0.05$, all realisations demonstrated an initial **epidemic** period over the first 6 weeks, as shown in Figs. 2.2a-c. After this period, the dynamics settled down to a quasi-equilibrium state characterised by fluctuations around a mean value that is the focus of our study – the following subsections are grouped by major findings of the model. The term *model variables* refers to the outputs $\langle S \rangle, \langle I \rangle, \langle R \rangle, \langle V_p \rangle, \langle N \rangle$ and $\langle V_s \rangle$.

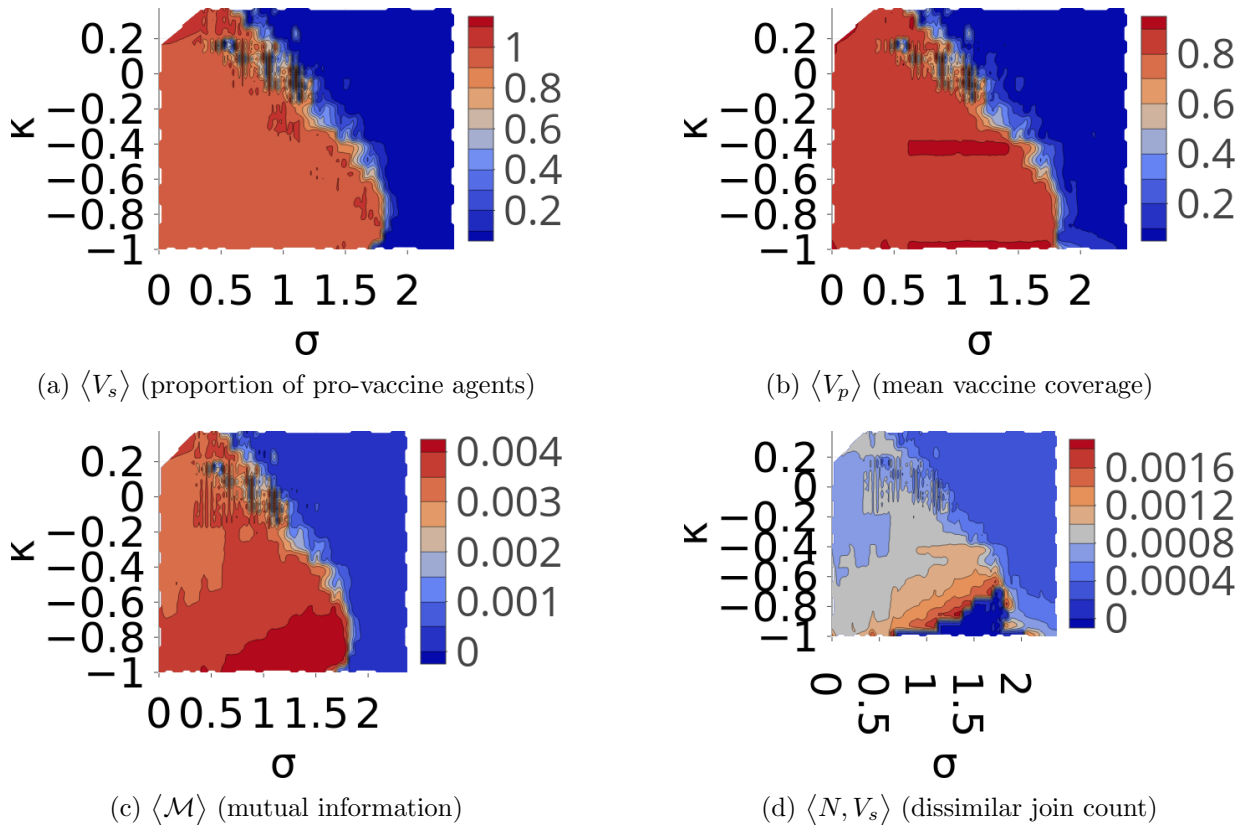


Figure 2.1: Contour plots of the region $(\sigma, \kappa) \in [0, 2.4] \times [-1, 0.2]$ of the parameter plane, capturing the transition dynamics of both the social and infection dynamics averaged over 20 realisations of each set of parameters; σ represents the strength of the social norm and κ the perceived vaccine risk.

2.3.1 Population vaccine immunity status can differ from aggregate vaccine opinion

Because only susceptible individuals are vaccinated and individuals cannot become ‘unvaccinated’ (but may change their opinion about the vaccine over their lifetime), the population-averaged vaccine opinion is not equal to the population-averaged vaccine immunity status, even at the quasi-equilibrium state. With no social pressure ($\sigma = 0$), a small increase in perceived vaccine risk $\kappa \rightarrow 0.03125$ pushes the system to endemic infection and anti-vaccine consensus (Fig. 2.2d) despite a high vaccination rate (Fig. 2.2g). Towards an explanation, if an agent n is newly birthed

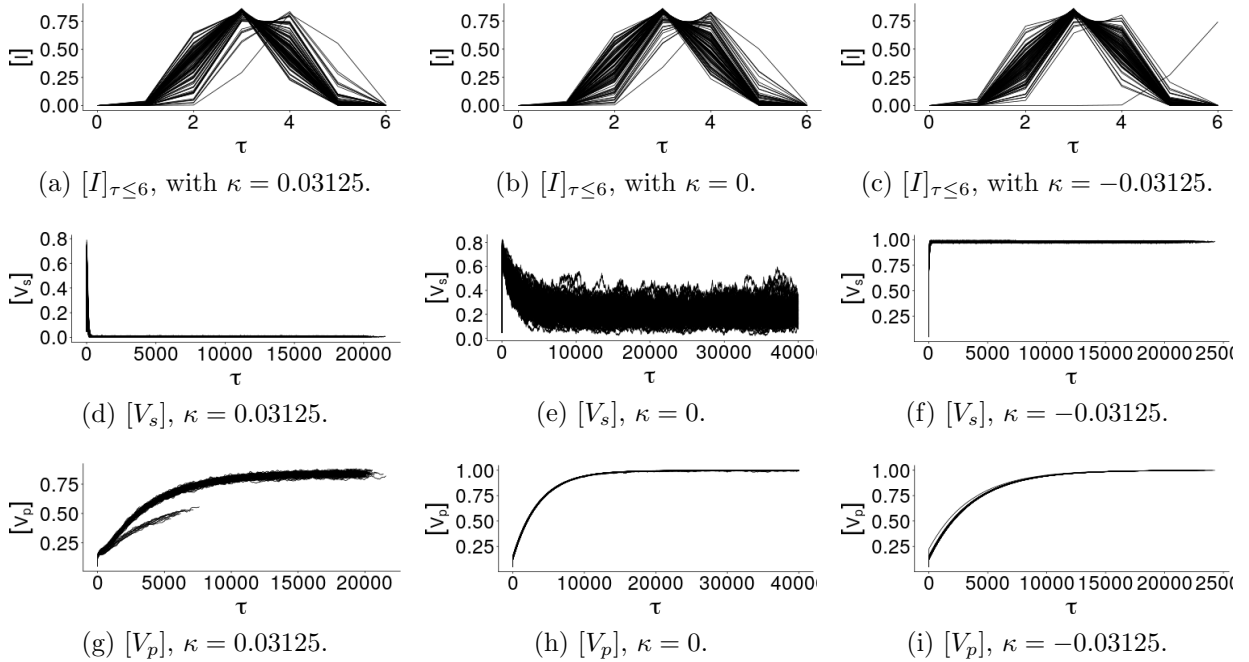


Figure 2.2: Time series demonstrating high sensitivity of the social dynamics to small changes (both positive and negative) in perceived vaccine risk κ when the strength of the social norm $\sigma = 0$. All panels show the results of 100 realisations of respective parameter combinations. τ represents the number of time steps (where a time step represents a week).

into this regime, the probability of having an infected neighbour vanishes ($\langle I_n \rangle \rightarrow 0$), so that

$$\mathbb{P}_n(N \rightarrow V_s) = \mathbb{P}_n(V_s \rightarrow N) \approx \frac{1}{2}, \quad (2.8)$$

similar to (1.1), with the agents' probability of being vaccinated over their lifetime as

$$0.05 + 0.95 \sum_{m=1}^{80} \frac{1}{2} \left(1 - \frac{1}{2}\right)^{m-1} \approx 1, \quad (2.9)$$

under the assumptions that the average agent with anti-vaccine opinion is almost certain to interact with a disagreeing neighbour. A similar calculation explains the phenomenon of high vaccination rates (Fig. 2.2h) despite a mixed consensus (Fig. 2.2e) when perceived vaccine risk becomes neutral ($\kappa \rightarrow 0$). A vaccine then perceived as beneficial ($\kappa \rightarrow -0.03125$) intuitively results in a high vaccination rate (Fig. 2.2i) and pro-vaccine consensus (Fig. 2.2f); in both these regimes, the infection survives only through case importation. In the absence of social norms and perceived vaccine risk, the population's aggregate vaccine opinion may not be a good indicator of

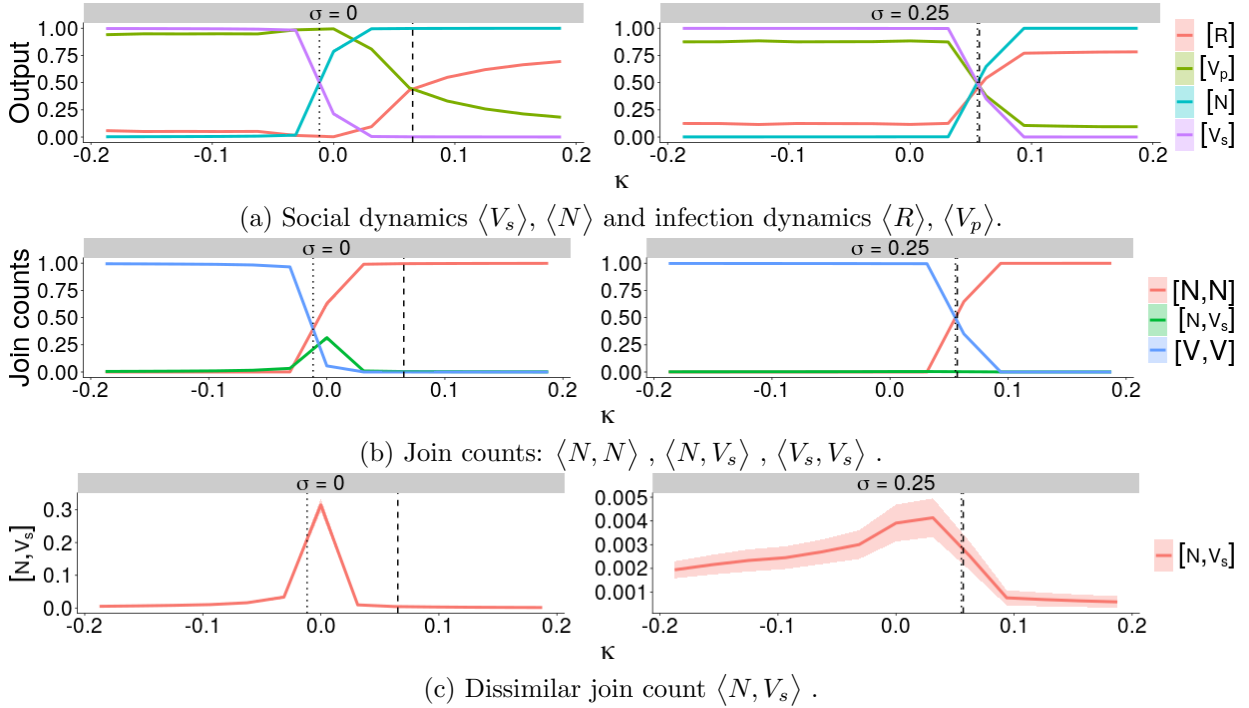


Figure 2.3: Trends of the equilibrium values of some EWS approaching the transitions of the social and infection dynamics K_s and K_p (marked by dashed and dotted vertical lines respectively), demonstrating the signals given by each tool with respect to the perceived vaccine risk κ . The intervals in each panel represent one standard deviation of the mean equilibrium value in each stochastic realisation of the model. Social norm $\sigma = 0$ for all panels on the left and $\sigma = 0.25$ for all panels on the right.

its vaccine immunity profile (and *vice versa*); (1.2) shows that the probability of changing opinion depends only on I_n when social norm $\sigma = 0$. In this region, the pattern of infection spread will be determined by the initial conditions of the infection dynamics; slight changes in perceived vaccine risk κ will push the network towards either of the consensuses, with minimal effect on the high vaccination rate (Fig. 2.2g). This phenomenon is shared by models V1 and V3, as shown in App. A.2.1.

2.3.2 EWS trends identify approaching transitions in both social and infection layers

The trends in both model dynamics and the proposed EWS are shown in Figs. 2.3 and 2.4, with $\sigma = 0$ (panels on the left) and $\sigma = 0.25$ (panels on the right). The dotted vertical line in

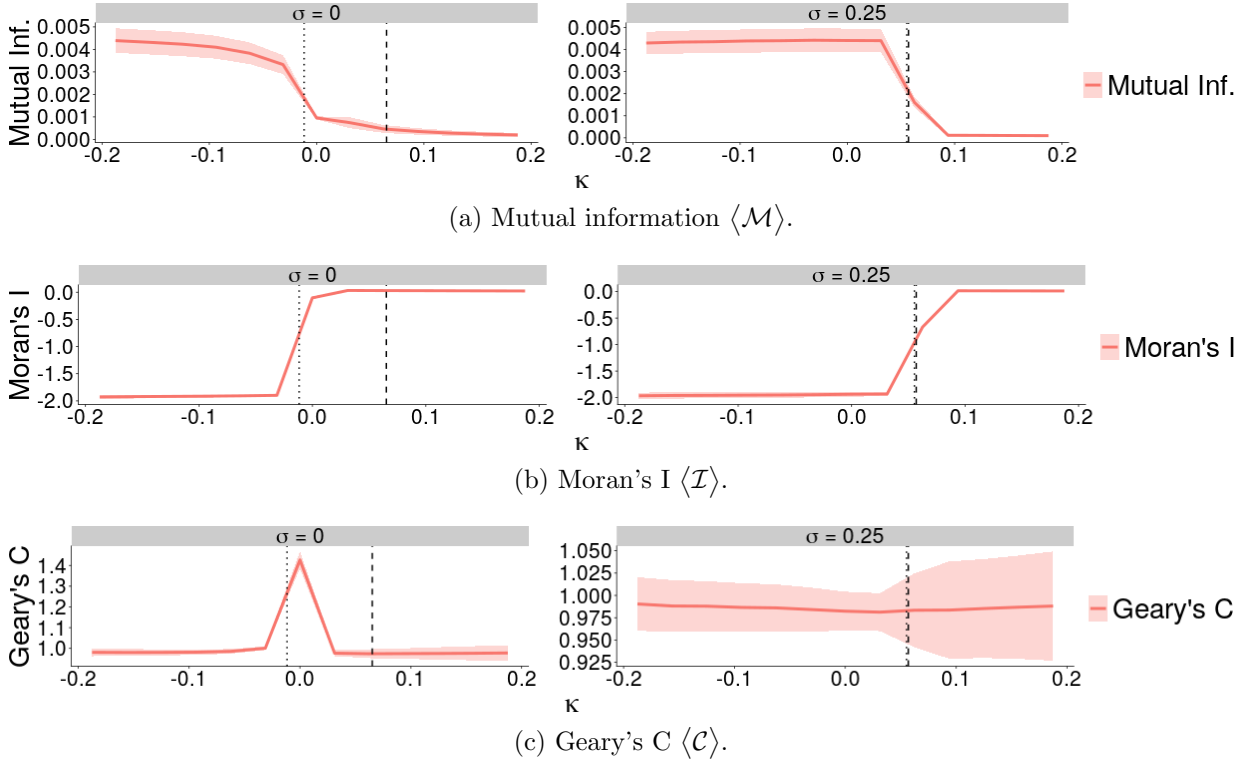


Figure 2.4: Trends of some of the EWS' equilibrium values approaching the transitions of the social and infection dynamics K_s and K_p (marked by dashed and dotted vertical lines respectively), demonstrating the signals given by each tool with respect to the perceived vaccine risk κ . The intervals in each panel represent one standard deviation of the mean equilibrium value in each stochastic realisation of the model. Social norm $\sigma = 0$ for all panels on the left and $\sigma = 0.25$ for all panels on the right.

all panels of Figs. 2.3 (as well as Fig. 2.4) represents the transition in the social dynamics K_s , defined as the smallest κ value at which $\langle V_s \rangle \approx \langle N \rangle$ (the mean number of pro-vaccine agents equals the number of anti-vaccine agents); the dashed vertical line represents the transition in the infection dynamics K_p , similarly defined as the earliest κ value at which $\langle R \rangle \approx \langle V_p \rangle$. Multiple infection and social transitions were found for some parameter combinations; these trends and the attendant behaviours of the EWS for models V1 and V3 can be seen in App. A.2.3. We also note that the equilibrium values of EWS and model variables were averaged over 15 – 20 realisations of all parameter combinations.

In Figs. 2.3 and 2.4, all EWS show recognisable trends preceding both transitions for both social norm values $\sigma = 0$ (panels on the left) and $\sigma = 0.25$ (panels on the right); for instance


$\langle N, N \rangle$ (Fig. 2.3b), $\langle N, V_s \rangle$ (Fig. 2.3c), $\langle \mathcal{C} \rangle$ (Fig. 2.4c) and $\langle \mathcal{I} \rangle$ (Fig. 2.4b) increase sharply preceding K_s with all but $\langle N, V_s \rangle$ approaching some maximum value preceding K_p , while $\langle V_s, V_s \rangle$ and $\langle \mathcal{M} \rangle$ sharply decrease and approach some minimum value before K_s and K_p respectively. Mutual information $\langle \mathcal{M} \rangle$ (Fig. 2.4a) in particular shows clear changes in trend well before the social transition K_s occurs. Though $\langle N, V_s \rangle$ (Fig. 2.3c) shows a similar rising-falling pattern for both $\sigma = 0, 0.25$, its maximum value with $\sigma = 0.25$ is much lower than that for $\sigma = 0$.

For $\sigma = 0.25$, the mean of the Geary’s C $\langle \mathcal{C} \rangle$ (Fig. 2.4c) shows almost no change, though its envelope broadens post-transition; we see this as a failure of the EWS (no forewarning given). Similar observations hold for model V1 (App. A.2.2), with the failure of the Geary’s C coefficient \mathcal{C} shown clearly in Figs. A.5l and A.6l. As stated in Sec. 2.2, the pre- and post-transition trends of the EWS do not generally resemble each other; asymmetry of the EWS about K_s can be seen Figs. 2.3, A.5 and A.6, showing that (in general) less of a warning is given (if any) when the κ -series is reversed. This is explicitly demonstrated in Figs. A.16 and A.15, where skewness γ_1 is used to quantify asymmetry of the trend of each EWS.


We can then say that all proposed EWS other than Geary’s C ($\langle \mathcal{C} \rangle$) give appreciable signals approaching K_s and K_p when $\sigma = 0.25$ (the right panels of Figs. 2.3 and 2.4). K_s precedes K_p (Fig. 2.3a), showing that a shift in consensus will always precede a crisis in vaccination coverage in this model. Also shown is a marked decrease in $K_p - K_s$ (the gap between the two transitions K_s and K_p , which we call the *intertransition distance*) as the social norm strengthens (for example, $\sigma \rightarrow 0.25$ in the right panels of Figs. 2.3 and 2.4). The generalisation of these trends to all tested values of σ is confirmed in Fig. 2.5a, where $K_p - K_s$ is everywhere positive, though the distance between K_s and K_p vanishes with increasing σ ; the inset of Fig. 2.5a shows the location of K_s (blue) and K_p (red) with respect to σ , so that $K_p - K_s$ (purple) gives the width of the area between the two curves in the inset graph at each σ . Other disparate models of the infection display largely similar concave decreases in the intertransition distance, suggesting that this behaviour arises generally from the model dynamics rather than in some specific subspace of the parameter space (see App. A.2.2).

2.3.3 Stronger social norms result in decreased lead distance for all EWS


The findings of the preceding subsection are intuitive, as vaccination depends more heavily on individual vaccination opinion than the number of infected neighbouring agents in (1.2), so that the opinion dynamics exert more influence than any feedback effect occurring in the infection dynamics. However, the vanishing intertransition distance $K_p - K_s$ presents a problem if we depend on predictions of K_s to enact interventions avoiding the collapse of the system to a non-vaccinated regime (i.e. avoiding K_p). For social norms of increasing strength, we can therefore look at the trend in the *lead distance* $K_s - \Pi_\Psi$, where Π_Ψ represents some κ value at which we can assert that a signal occurs in some sequence Ψ of κ values; since we’ve established that K_s



(a) Demonstration of the shrinking intertransition distance $K_p - K_s$ (purple), with the inset graph showing the estimated locations of K_s (blue) and K_p (red).



(b) Under the Lanzante change point test, the lead distance of each EWS ($\text{Lan}_\sigma\{\text{WS}\}$) varies substantially with the strength of the social norm σ ; trends corresponding to each EWS are represented by the different colours in the legend; the bar chart on the right gives the percentage of σ values for which each individual EWS gave the maximum lead distance.



(c) Panel showing the variance of lead distances ($\text{Lan}_\sigma\{\text{S/I}\}$) of the model variables with social norm σ , with the bar chart on the right giving the percentage of times each model variable gave the maximum lead distance of all variables (over all values of sigma).

Figure 2.5: Figure showing the decreasing trend of the intertransition gap $K_p - K_s$, as well as the variation of the lead distances of EWS and model variables obtained by applying the Lanzante change-point test to their respective κ -series.

precedes K_p everywhere, then necessarily any warning of a social transition also warns of the following infection transition, so the quantity $K_p - \Pi_\Psi$ is not discussed here.

One way for us to quantify this lead distance is to use a [change point detection test](#) to find κ values at which the two classes EWS (denoted **WS**) and model variables (denoted **S/I**) give signals (i.e. display statistically significant changes in trend/behaviour). Here, the Lanzante test [319] from the **trend** [454] package in R is applied to various sequences of equilibrium κ values to find change points of EWS ($\text{Lan}_\sigma\{\text{WS}\}$, Fig. 2.5b) and [model variables](#) ($\text{Lan}_\sigma\{\text{S/I}\}$, Fig. 2.5c) respectively. (In other words, we computed the change test on the mean of all stochastic realisations at each parameter value combination, rather than computing the change test on each individual time series and then taking the average.) Further discussion of the method of application of this and other change point detection tests to series of κ values can be found in App. A.2.4. *Failure* of a warning signal or model variable Ψ occurs when the warning given comes *after* the social transition, so that $K_s < \text{Lan}_\sigma\{\langle\Psi\rangle\}$, where $\text{Lan}_\sigma\{\langle\Psi\rangle\}$ represents the change point obtained from the Lanzante test.

We denote the lead distance $\text{Lead}_\sigma[\langle\Psi\rangle]$ of some EWS Ψ at social norm σ with (Lanzante test) change points as

$$\text{Lead}_\sigma[\langle\Psi\rangle] = K_M - \text{Lan}_\sigma\{\langle\Psi\rangle\}; \quad (2.10)$$

looking at the trend in the lead distances $\text{Lead}_\sigma[\langle\text{WS}\rangle]$ (Fig. 2.5b), the positivity of some curves shows that some of the proposed EWS do indeed give early warnings of coming transitions (largely for $\sigma \leq 1.875$). Failures of $\langle N, V_s \rangle$ and $\langle \mathcal{M} \rangle$ occur in the range $1.875 \leq \sigma \leq 2.125$, while all other tests give valid warning signals everywhere $\sigma < 2.5$; model variables $\langle R \rangle$ and $\langle V_p \rangle$ fail in the range $1.65 \leq \sigma \leq 2.125$ (Fig. 2.5c). The failure of all the tests after $\sigma = 2.5$ likely results from insufficient length of the EWS' κ -series; the inset of Fig. 2.5a shows that $K_s \rightarrow -1$ as σ increases. Figure 2.5c is largely similar, showing failure of all the signals around $\sigma = 2.5$ (as in Fig. 2.5b). In line with our focus on social dynamics as a predictor, $\langle V_s \rangle$ appears to be the best performing signal of all the model variables; as was reasonably expected, $\langle I \rangle$ appears to perform badly, since its role as a transitory compartment means that it never ‘‘gathers’’ sufficient agents over the course of each realisation to give a true indication of the state of the system (other than indicating the presence or absence of endemic infection).

Since the perceived risk of vaccination $\kappa \geq -1$ in this study, our method of detecting the change point will not accurately predict a change point Π_* close to -1 . Since no one warning signal gives the highest lead distance for any large contiguous range of σ values, there is unfortunately no single objective way to choose a ‘‘strongest’’ signal; they are all suitable tools to predict coming crises in aggregate opinion and vaccination dynamics. However, it is worth noting that mutual information $\langle \mathcal{M} \rangle$ and the dissimilar join count $\langle N, V_s \rangle$ perform better than the other indicators; Fig. 2.5b shows that mutual information $\langle \mathcal{M} \rangle$ give the largest lead distance measured for 45% of tested social norm σ values, and $\langle V_s, V_s \rangle$ and $\langle N, V_s \rangle$ both give the largest lead distance for 32% σ values (multiple EWS showed an identical lead distance for some values of σ). Lead distances of all EWS for all three models under other various change point tests are discussed in App. A.2.4.

2.3.4 There's no single best EWS per σ value

We can compare the EWS by quantifying their tendencies to give maximum and minimum lead distances. This is shown in Fig. 2.6, where green bars represent the ratio of tested σ values for which the EWS gave the largest lead distance of all the EWS and red bars represent the ratio for which the EWS gave the lowest lead distance.

Mutual information $\langle \mathcal{M} \rangle$ gave the largest lead distance for 50% of σ values (Fig. 2.6a), with the next best EWS (in order) being the dissimilar join count $\langle N, V_s \rangle$ (with 41% maxima of the EWS) and the pro-vaccine similar join count $\langle V_s, V_s \rangle$ (32% maxima of the EWS). Both gave the minimum lead distances for 14% of σ values, the lowest ratios of the lot. Geary's C $\langle \mathcal{C} \rangle$ was also one of the worst of the EWS, giving the lowest lead distance for 36% of σ values.

Also of interest is whether the performance of any EWS depended on the range of σ values; dependence would be indicated by a discernible pattern in the colours of the squares in any row of Fig. 2.6b, which categorises the lead distance of each EWS by σ value. The length of an EWS' green bar (maximum lead distances) in Fig. 2.6a represents the proportion of green tiles in the EWS' row of Fig. 2.6b; the same relationship hold between the red bars of Fig. 2.6a and red tiles in Fig. 2.6b. There is no such discernible pattern of either green (EWS maxima) or red (EWS minima) squares for any EWS. Grey squares represent neutral values (neither maximum nor minimum) and yellow squares signify equal warning among all EWS. The seeming lack of pattern in Fig. 2.6b shows that any notion of *best* performance of any EWS cannot be described by a specific subset of σ values. Similar analysis of models V1 and V2 under various change point tests can be found in App. A.2.7, as well as further argument for why the anti-vaccine similar join count $\langle N, N \rangle$ and Geary's C $\langle \mathcal{C} \rangle$ statistics should not be used for the models investigated.

2.3.5 EWS can provide better forewarning than trends in model variables


One final question is whether the proposed EWS (mutual information, Moran's I, Geary's C, join counts) give earlier warnings than simply monitoring trends in model variables (such as the number of infections or pro-vaccine agents, using a change point test for prediction in both cases). There are many ways to quantify this, including **maximin** comparison (finding the larger of minimum values of classes **WS** and **S/I**) and **maximax** comparison (finding the larger of the maxima of each class) at each σ .

To compare the minima of the EWS and model variable lead distances, we define χ_{\min}^{Lan} as


$$\chi_{\min}^{\text{Lan}} = \min_{\sigma} (\text{Lead}_{\sigma}[\langle \text{WS} \rangle]) - \min_{\sigma} (\text{Lead}_{\sigma}[\langle \text{S/I} \rangle]) , \quad (2.11)$$

and we specify a tolerance $\varepsilon_{\min}^{\text{Lan}}$ to be 1% of the total range of χ_{\min}^{Lan}

$$\varepsilon_{\min}^{\text{Lan}} = \left| \frac{\max(\chi_{\min}^{\text{Lan}}) - \min(\chi_{\min}^{\text{Lan}})}{100} \right|. \quad (2.12)$$



(a) The proportion of tested σ values for which each respective EWS gave the largest (green bar) and smallest (red bar) lead distance.



(b) Grid showing the relative performance of each EWS with the Lanzante change point test. Green tiles denote the σ values for which the EWS gave the highest lead distance, red tiles represent the smallest lead distance, grey tiles represent lead distances that are neither maxima nor minima and yellow shows where all EWS gave the same lead distances. Black tiles represent failed warnings (negative lead distances) and white tiles represent undefined values.

Figure 2.6: Mutual information $\langle \mathcal{M} \rangle$ is the best EWS in (a), giving the largest lead distance of all EWS for 50% of σ values; the dissimilar join count is second best 41%, with both giving the smallest lead distances for only 14% of σ values. (b) shows that there is no pattern in EWS performance.

Hence, if $\chi_{\min}^{\text{Lan}} > \varepsilon_{\min}^{\text{Lan}}$, then the EWS is outperforming simple monitoring of trends (model variables).

The blue curve in Fig. 2.7 compares the minima of the EWS (WS) and model variable (S/I) classes (maximin comparison), showing the σ values for which the worst-performing (least lead




Figure 2.7: Graph of the trends of χ_{\min}^{Lan} (blue) and χ_{\max}^{Lan} (red) with respect to the value of the social norm σ , allowing us to do maximin and maximax comparisons of the two classes of warning signals (WS and model variables S/I). The green-shaded region shows where $\chi_*^{\text{Lan}} > 0$ and the red-shaded region shows where $\chi_*^{\text{Lan}} < 0$. The inset table shows the percentage of σ values for which $\chi_*^{\text{Lan}} > \varepsilon_*^{\text{Lan}}$ (*pos*: EWS works better), $|\chi_*^{\text{Lan}}| \leq \varepsilon_*^{\text{Lan}}$ (*zero*: both approaches work equally well) and $\chi_*^{\text{Lan}} < -\varepsilon_*^{\text{Lan}}$ (*neg*: monitoring simple trends works better). **(blue curve, row 1 of inset table)** Positive values (green-shaded region) of χ_{\min}^{Lan} occur at the σ (social norm) values where the worst-performing (least lead distance) EWS still gives higher lead distance than the worst-performing model variable. **(red curve, row 2 of inset table)** Similar to above, positive values of χ_{\max}^{Lan} occur (in the red-shaded region) when the EWS perform absolutely better than the model variables.

distance) EWS $\min_{\sigma} (\text{Lead}_{\sigma}[\langle \text{WS} \rangle])$ is either better, *equal* or worse than the worst-performing model variable $\min_{\sigma} (\text{Lead}_{\sigma}[\langle \text{S/I} \rangle])$. Points in the green-shaded region represent σ values where the EWS' performance is at worst still better than that of the model variables.

The EWS outperformed simple monitoring of trends in variables for 45.4% of the tested σ values; maximin comparison shows that EWS are at worst still better predictors than the model variables for a large number of σ values, with the two classes performing equally badly in 18.2% of the σ values. Performance of the EWS and the model variables in this test was considered comparable or *equal* if the difference between the two minimum lead distances fell under the tolerance ϵ_m , so that $|\chi_{\min}^{\text{Lan}}| \leq \varepsilon_{\min}^{\text{Lan}}$; performance was equal for 18.2% of tested σ values, showing that the added computation of the EWS does not always yield a benefit. Otherwise, the points and portion of the blue curve falling in the red-shaded region of Fig. 2.7 represents values of σ where the model variables outperformed the EWS (that is, the minimum lead distance of the model variables exceeded the minimum lead distance of the EWS); this occurred for 36.4% of tested σ values.

The second part of the comparison (Fig. 2.7, red curve) is between the maxima of the lead

distances; as above, we define the comparison variable χ_{\min}^{Lan} and tolerance $\varepsilon_{\max}^{\text{Lan}}$ as

$$\chi_{\max}^{\text{Lan}} = \max_{\sigma} (\text{Lead}_{\sigma}[\langle \text{WS} \rangle]) - \max_{\sigma} (\text{Lead}_{\sigma}[\langle \text{S/I} \rangle]) , \quad \varepsilon_{\max}^{\text{Lan}} = \left| \frac{\max(\chi_{\max}^{\text{Lan}}) - \min(\chi_{\max}^{\text{Lan}})}{100} \right| . \quad (2.13)$$

The green-shaded portion of Fig. 2.7 also shows the σ values where the EWS outperformed the model variables here, in that the maximum lead distance given by the EWS exceed that given by the model variables ($\chi_{\max}^{\text{Lan}} > \varepsilon_{\max}^{\text{Lan}}$); points falling within the red-shaded area of Fig. 2.7 show for which σ values the model variables outperform the EWS. From the second row of the inset table in Fig. 2.7, the two maxima are considered equal ($|\chi_{\max}^{\text{Lan}}| < \varepsilon_{\max}^{\text{Lan}}$) for 63.6% of tested σ values, while the EWS outperformed the model variables ($\chi_{\max}^{\text{Lan}} > \varepsilon_{\max}^{\text{Lan}}$) for only 13.7% of σ values.

This shows that the EWS' lead distances are at least equal to those of the model variables for around 63.6% of σ values and are absolutely larger for 77.3% of σ values, demonstrating that though monitoring the model variables (both social and infection dynamics) is itself valuable, the EWS offer better performance (using the Lanzante change point test). In both (blue and red) curves of Fig. 2.7, there is no apparent pattern to the positivity/negativity of χ_{\min}^{Lan} and χ_{\max}^{Lan} . These comparisons are given for other tests and models in App. A.2.4.

2.4 Discussion

Here we studied a range of [early warning signals](#) for critical transitions in a coupled behaviour-infection model of paediatric infectious diseases. We compared the indicators to one another and the approach of simply monitoring trends in model variables. We found that the performance of the indicators was variable depending on model parameters, but the mutual information statistic and the dissimilar join count showed consistently high pre-transition lead distances over various strengths of the social norm σ , many times giving the highest lead distances of all the EWS. Through maximin and maximax comparisons, we found that using EWS provide more advance warning than simply monitoring trends in model variables in a clear majority of cases.

We note that join counts have the additional advantage of easy computability, since they require only counting pairs of a given type. This contrasts with other more computationally intensive indicators (such as autocorrelation) which require making decisions about whether to study lag-1 or higher order lags, as well as choosing parameter values governing computation of residuals. Moran's I was also shown to predict the approach of transitions, although perhaps this finding is trivial considering that it is a linear combination of similar and dissimilar join counts. Its predictive power was not as strong as many of the other indicators such as join counts and mutual information, hence the added complexity of its calculation may not justify its use. Potential downfalls of the mutual information statistic include its computational complexity and

the availability of a suitable data set pairing the personal health of each agent with their social activity.

Under an increase in the perceived vaccine risk κ the model exhibited an earlier critical transition in the social layer than in the infection layer (especially for low social norms σ). We also showed that a population may have relatively high vaccine coverage despite a low pro-vaccine opinion. This discrepancy between social and infection dynamics is due to the paediatric infectious disease natural history we assumed in the model. Unlike influenza, where revaccination must occur seasonally, an individual who receives a sufficient number of measles or chickenpox vaccine doses generally has lifelong immunity and therefore the opinion towards the vaccine can decline well before the level of vaccine immunity does. (Individuals can change their opinion but never become ‘unvaccinated’.) This asymmetry between population opinion and vaccine immunity status in the model is reflected in several real-world populations. For instance, a recent survey in France [373] indicates that only two in three residents agreed that vaccines are safe [536] while the vaccine coverage rates for measles and HepB3 (the third dose of the Hepatitis B vaccine) were both 90% amongst children [37, 36], with a 96% rate of DPT inoculation [35]. The implication of this asymmetry is that monitoring social media networks for changes in opinion using early warning signals like mutual information might provide advance warning of outbreak hot-spots.

The distance between the change point in the EWS indicators and the critical transition in the social dynamics decreases as the social norm grows stronger, as does the distance between the transitions in social and infection dynamics of the model. Given the relative scale of the social norm and perceived vaccine risk parameter values used, stronger social norms decrease the time interval between birth and vaccination decision (the vaccination rate converges to its equilibrium value in fewer time steps than in other regimes); feedback between this and the infection incidence in the network (which affects the number of infected neighbours in each agent’s neighbourhood) alters the probability function controlling the vaccination decisions, effecting faster alignment of majority opinion and vaccination coverage.

This study only lays the foundation for the investigation of spatial EWS for such coupled-behaviour systems. There is much work to be done and many questions answered before they can be meaningfully applied to empirical data. For instance, we assumed that the network was static. This simplifying assumption could be relaxed in future work by using an evolving social dynamic in which agents are allowed to form or break links with new agents based on their node degree [395] or vaccine opinion and associated social pressures [187, 394], mimicking assortative/disassortative link formation on social networks.

Likewise, triadic closures could be allowed to capture clustering [263, 185]. The rate of interaction between agents was also assumed fixed in our model (relative to the speed of other dynamics, such as the birth/death interval and the length of illness). A valid extension of the model would be a variable rate of communication between agents, since the rate of communication has been shown to influence the rate and efficiency of opinion formation [434]. A further avenue of research

would explore interventions to turn populations away from critical transitions. This could answer research questions such as: *How far in advance must we act to prevent a social or infection critical transition and does this change our interpretation of the EWS? If we assume that any of the EWS can be used for monitoring, how would this change in trend alter the reliability of the EWS?* Finally, the phenomenon of vaccine hesitancy (in contrast to clear anti-vaccine opinion) is both widespread [496, 344, 166] and dangerous [506]; any substantial change to social dynamics effected by the inclusion of vaccine hesitancy to this binary model may alter the responses of the EWS since they all depend on the balance of pro-/anti-vaccine sentiment in the neighbourhood of each agent. This robustness of the EWS indicators the inclusion of vaccine hesitancy would be of particular interest.

Many researchers and public health bodies are drawing attention to global resurgences of vaccine-preventable illness and speaking to the vast efforts and multiple approaches taken to mitigating outbreaks. A few of these approaches have focused on human behaviour and opinion dynamics, either by directly tracking aggregate vaccine opinion, or monitoring alerts and media reports. Our work demonstrates the potential uses of early warning systems of critical transitions in preventative epidemiology. In particular, our work provides proof-of-concept for the idea of monitoring social networks for early warning signals of both social and epidemiological shifts, and also suggests several EWS indicators that might work well for this purpose.

Chapter 3

Community structure predicts social shifts

Abstract

Sudden shifts in population health and vaccination rates occur as the underlying dynamics of some epidemiological models go through a critical point; literature shows that this phenomenon is sometimes foreshadowed by early warning signals (EWS). We investigate different structural measures of a network as candidate EWS of infectious disease outbreaks and significant changes in aggregate vaccine sentiment. We construct a multiplex disease model coupling infectious disease spread and social contact dynamics. We find that the number and mean size of echo chambers of pro- and anti-vaccine agents predict critical transitions in the infection dynamics, as with the communities on the social network. The graph modularity measure also gives early warnings, though the global clustering coefficient shows no significant pre-outbreak changes. Four change-point tests applied to series of the EWS show decreasing efficacy of the measures as social norms strengthen. This shows that many graph-theoretic measures of social network connectivity can be used to predict approaching critical changes in vaccine uptake and aggregate health in various populations, thereby providing valuable tools for improving public health.

3.1 Introduction

Low vaccine rates stemming from [vaccine](#) refusal [443] result in outbreaks of vaccine-preventable diseases in some populations [105, 469]. The high costs of intervention and treatment incurred by public health systems [231, 350] motivate us to find tools warning of [epidemics](#). The connection between social network activity and health issues in populations has long been exploited

by researchers [461, 88, 115, 116, 181], especially relating to disease spread [478, 584, 17, 208], with the assertion that firm understanding of social network structure is important to the implementation of effective policy interventions [244]. For example, vaccination decisions sometimes depend on communication and information diffusion in media [574, 533, 506]; as such, patterns of communication in social networks yield warning signals such as increased spatial autocorrelation [163, 274, 420].

Through **assortative mixing**, densely connected groups of members are formed in a social network with sparse connectivity between said groups; these groups are called *communities*, and greatly influence any dynamic process on the network [224, 123, 317, 33, 58, 218, 376, 370, 116, 115, 203, 95]. Specifically, studying the formation and growth of **communities** on social networks allows for the discovery of non-obvious interrelationships [317, 572, 530]. Strongly connected components on social networks have been used as a proxy for community structure [328] as well as other structures and definitions [465, 315, 134], with different benchmarks used to compare the efficacy of different methods of detection [224, 126, 140, 316, 314, 315, 419].

Communities where every member shares the same opinion are called *opinion clusters* [370, 478, 244] or *opinion based communities* [106] and are a fixture of online social networks, a direct result of assortative mixing [106, 61, 62, 154]. Different methods of propagation of opinion have been studied in the literature: for example, neighbourhood sampling [554, 100], summation and averaging [251, 207, 99], estimation [349] and population-level interaction [211, 553]. In communication models where sentiment change is driven primarily by exposure to news sources and contrasting views from neighbours, communities can support the reinforcement of sentiments already held [537]. For example, in a model where each node has a numerical strength of opinion ranging from 0 to 1 with nodes only able to influence each other if the difference between their opinions falls below a threshold, higher values of this threshold exacerbate the formation of large opinion clusters [244, 370].

Echo chambers, described as well connected groups of people promoting and reinforcing the same bias [154, 463, 155, 490], have recently come under media scrutiny since these groups can facilitate vaccine scares [490], support political candidates [106], lead to skewed evaluation of objective fact and decreased accuracy of opinion [337]. Furthermore, some studies indicate that in some cases these homogeneous subnetworks may reinforce bias [518, 488], in some part due to avoidance of cognitive dissonance [307, 219]. Given that interaction between dislike agents in a network can sometimes correct false beliefs [497, 240], these echo chambers may be seen as drivers of polarisation [519]. This is especially since much anti-vaccine content is shared without thought of its veracity [16, 294, 308, 383, 64, 65], with Facebook anti-vaccine groups serving the dual purpose of opinion-reinforcing echo chamber and “fake news” source in a time where a large number of people draw on social media sites for their health information [111, 84, 277, 462, 303, 489, 63].

In the same vein, much work has focused on the modularity of social networks. **Modularity** is a graph theoretic measure of the segregation of a graph [391]; a high degree of modularity

may indicate increasing segregation of a network into clusters [128], with other work showing that modularity is not a “direct measure of polarisation” [239]. With some governing dynamic, modularity is a “essentially rooted at the stability of its corresponding social system”, with stable networks containing one large community and unstable networks showing modularity driven by polarisation [569]. The [global clustering coefficient](#) (GCC) works in a similar way; by describing the number of triangles in the network, it is an important measure of graph structure [351]. A high clustering coefficient is accompanied by the small world phenomenon of short average inter-node distance [525] which facilitates efficient information spread on social networks through redundancy [103, 102, 338]. This clustering also facilitates mutual communication between nodes, where effective communication between disagreeing persons can lead to change of opinion [106].

The occurrence of opinion change within opinion-based communities has seen much attention in political studies, with some work asserting the ability of the ‘wisdom of the crowd’ to overcome bias [54], while other work shows that group phenomena reinforce the opinions held [521, 337, 380]. Others have argued that a commonly held belief within a community can still become more accurate even as the homogeneity of the group increases [53, 273]. This leads to our interest in the rate of opinion change in the network as yet another potential indicator of dynamical regime change.

As mentioned in Ch. 2, systems moving from one polarised state to another undergo [phase transition](#) through a sole [critical point](#) [385]. Called *critical transitions*, they sometimes result in a demonstration of characteristic system behaviours such as *critical slowing down* [484, 415]. These events give us easily recognisable ‘hints’ of approaching transitions called *early warning signals* [81, 484, 485]. Here, we apply Ch. 2’s methodology to a new set of prospective EWS.

We show that trends in all of the measurements described above ([modularity](#), [global clustering coefficient](#), census and sizes of communities and [echo chambers](#)) provide early warning signals of epidemic and vaccine crisis events for a coupled disease-behaviour model of childhood disease. We use a binary vaccine opinion dynamic and an $SIRV_p$ disease process occurring on a random network to model a childhood infectious disease. By quantifying and comparing their performance, we find that trends in the sizes of anti-vaccine communities were the best-performing signals, with the modularity and clustering coefficients of communities also performing well. All in all, we verify that changes to fundamental graph structure driven solely by opinion dynamics are good predictors of disease events and aggregate sentiment towards vaccination.

This chapter is organised as follows: in Sec. 3.2, we will describe the disease-behaviour model used in the study and a description of the warning signals used. Section 3.3 will show the change in trends of the signals with respect to perceived risk of adverse vaccine effects and the social pressure of an injunctive norm, as well as a quantification of the warning provided by each measure through the use of the Standard Normal Homogeneity change point test ([SNHT](#)). We will elaborate on the limitations of the measures used and the implication of the results in the Sec. 3.4.

3.2 Methods

3.2.1 The model

For simulation, we use an [agent-based model](#) identical to that described in Sec. 1.3. A perfectly effective [vaccine](#) with no effect on mortality is immediately available to susceptible agents upon gaining pro-vaccine status $S \rightarrow V_p$. As shown in Fig. 1.7a, the disease follows the $SIRV_p$ model; each agent physically interacts with all their neighbours per time step. If a susceptible agent becomes infected $S \rightarrow I$ (with probability $1 - (1 - p)^{I_n}$, where I_n represents the number of infected neighbours of the agent n and $p = 0.2$ the probability of infection), they become both ill and infectious for $\ell = 2$ time steps (each time step represents a week); recovery $I \rightarrow R$ and vaccination $S \rightarrow V_p$ are both permanent. Injunctive social norms $0 \leq \sigma \leq 2$ and a perceived vaccine risk of $-0.4 \leq \kappa \leq 0.2$ represent peer pressure and adverse vaccine effects respectively [413].

Figure 1.7 features a binary social (opinion) dynamic, where agents demonstrate either pro-vaccine V_s or anti-vaccine N sentiment. Change of sentiment occurs through imitation of neighbours, where a randomly chosen neighbour is sampled each time step; an [effective interaction](#) with a disagreeing neighbour (with different sentiment than the agent sampling) prompts a reevaluation and change of sentiment with probability $\mathbb{P}_n(N \rightarrow V_s)$ for anti-vaccine agents and $\mathbb{P}_n(V_s \rightarrow N)$ for pro-vaccine agents. Any susceptible agent that adopts pro-vaccine sentiment ($N \rightarrow V_s$) is immediately vaccinated ($S \rightarrow V_p$); for agent n , these changes in sentiment depend on the perceived vaccine risk κ and neighbours I_n :

$$\begin{aligned}\mathbb{P}_n(N \rightarrow V_s) &:= \frac{1}{1 + \exp\left(-U_n^{N \rightarrow V_s}\right)}, \\ \mathbb{P}_n(V_s \rightarrow N) &:= \frac{1}{1 + \exp\left(-U_n^{V_s \rightarrow N}\right)}.\end{aligned}\tag{3.1}$$

Indices $U_n^{N \rightarrow V_s}$ and $U_n^{V_s \rightarrow N}$ in (3.1) are utility functions defined as

$$\begin{aligned}U_n^{N \rightarrow V_s} &:= -\sigma \left(\frac{d_n^N - d_n^{V_s}}{Q_n} \right) - (\kappa - I_n), \\ U_n^{V_s \rightarrow N} &:= -\sigma \left(\frac{d_n^{V_s} - d_n^N}{Q_n} \right) + (\kappa - I_n),\end{aligned}\tag{3.2}$$

where d_n^* represents the number of neighbours of n with sentiment $*$, with d_n representing the total number of neighbours.

3.2.2 Early warning signals

Four of the six **EWS** explored in this paper are related to the detection of community structure in networks; the **global clustering coefficient**, **echo chambers**, opinion-based **communities** and modularity score are all important tools that enable biological modelling [391]. The topological phenomenon of *community* refers not to a single central construct, but rather a general notion of variation in connection density. Communities in social networks are vaguely defined in the literature as groups with the following basic property: members of the community are more connected to each other than with non-members [465]. Vagueness in science usually leads to artistic licence and the specific treatment of context and purpose; by the above definition, communities can then be alternately conceptualised in different areas and studies as modules, clusters, groups and so on [550]. Such refinements then lead to tighter and more technical definitions.

Here, we conceptualise topological **communities** as (connected) components on the network. Components in undirected networks are maximal disjoint groups of agents such that there is a path between every pair of agents in the group. [539, 174]. Even more specifically, the concept of a *giant connected component* (**GCC**) describes a component that contains a “significant fraction of all the nodes” [174, 141]. The role of components (both non-giant and giant) in spreading processes can be conceptualised as such: where some infection can be spread from person to person, GCCs are formed by historical person-to-person contacts (as opposed to current contacts) [174, 51]. Practical indirect analysis and exploitation of this component structure in policy design and epidemiological intervention is facilitated through *contact tracing* [118, 117, 141].

Recent political upheaval has thrust the phenomenon of the *echo chamber* into social consciousness and modern parlance, with many lay articles arguing for and against their existence [311, 473, 167, 232, 87, 145, 470, 235, 398, 202, 467, 235, 259, 491, 512, 439], though many articles do not directly define the concept before exploiting it [77, 124, 578, 223]. There exist different definitions of echo chambers in literature; also called *tribes* [583, 582, 148] (sometimes sardonically in popular media), they can be described either as a community where at least some percentage of the members hold a particular sentiment [40, 127], or else a subset of community members overwhelmingly likely to restrict their neighbourhood communication to contacts with shared opinion [243, 397]. As such, **echo chambers** are usually conceptualised as closed subsystems of social networks containing members with a single orientation [537, 40], and are therefore considered synonymous with homophily and strongly associated with the quick spread of misinformation [168, 153, 361], polarisation and the insulation and reinforcement of belief despite their veracity [124, 537, 583]. One example is the concept of the ‘*filter bubble*’, detailing the biased filtering of information unfortunately created by personalisation algorithms [86, 198, 427, 171].

Given the described ubiquity on social networks and their importance in information diffusion (and therefore decision making) [168, 153, 361, 151], we test whether observations of the size and number of communities Z_* and echo chambers J_* give early warnings of approaching vaccine scares and crises, as well as falls in vaccination rates. We retrieve the echo chambers J_* by first




Figure 3.1: Community structure of pro-vaccine agents on a small network; red nodes represent anti-vaccine agents (N) and green nodes represent pro-vaccine (V_s) agents. The four grey convex hulls (one of which is a single node) represent pro-vaccine components and the yellow convex hull highlights the lone pro-vaccine echo chamber.

listing the agents in the network with the desired vaccine opinion; the members of the network are then sorted into two primary classes. The *peripheral* members maintain links with at least one disagreeing neighbour and *core* members only maintain links with agreeing neighbours (these are direct analogies of the boundary and interior of a topological space respectively). We then characterise echo chambers as communities of core members. This can be seen in Fig. 3.1, which presents a small dense network of pro- and anti-vaccine agents. Connected groups of pro-vaccine agents are indicated as pro-vaccine communities by the four grey convex hulls, one of which is a single node; the nodes captured in the grey areas are the peripheral members mentioned above. Within the largest community, there are only two agents which have no disagreeing neighbours; these are the core members. Since they are connected to each other, they form a single subcommunity which we refer to as an echo chamber.

Clustering has been shown to greatly facilitate the spread of such phenomena as political extremism [151] and infectious disease [479]. Specifically, clustering refers to the propensity of connection between two persons if they have a mutual friend [389]. As an indicator of this organisation, the

global clustering coefficient indicates the prevalence of clusters, dense highly-connected groups of nodes in a graph [412]. This done by finding the density of *triplets* on the network. An *open triplet* is a group of three nodes connected by 2 edges, while a *closed triplet* is a group of three nodes joined by three unique edges (also called *triangles* for this reason). The *global clustering coefficient* (C_*) is then calculated as

$$C_* = \frac{\sum \mathcal{T}_\Delta}{\sum \mathcal{T}_\Lambda}, \quad (3.3)$$

where $\sum \mathcal{T}_\Delta$ represents the number of closed triplets and $\sum \mathcal{T}_\Lambda$ represents the number of open triplets [551]. The value of the coefficient is bounded $0 \leq C_* \leq 1$. In metaphor with percolation theory literature, the clustering coefficient of subnetworks formed by holders of either sentiment could potentially act as an indicator of organisation that enhances the reinforcement of sentiment and the effect of any social norm.

The *modularity* measure is similar to the global clustering coefficient as a measure of organisation; highly modular networks possess many *modules*, which feature dense interconnectivity between nodes similar in some way and sparse connectivity between dislike nodes. Specifically, this measures the correlation between the probability of connection of two nodes and their membership of the same module [71]. The modularity score is calculated as follows [393]; let an undirected network be divided into two disjoint groups Λ and Ξ , with each node n given the score

$$s_i = \begin{cases} 1 & j \in \Lambda \\ -1 & j \in \Xi \end{cases}, \quad (3.4)$$

with an adjacency matrix A of the network, so that A_{ij} gives the number of edges between nodes i and j . Let k_* represent the degree of node $*$, so that

$$m = \frac{1}{2} \sum_i k_i \quad (3.5)$$

gives the number of edges in the network and the expected number of edges between nodes i and j is

$$\frac{k_i k_j}{2m}. \quad (3.6)$$

The network modularity (Q) is then given as

$$Q = \frac{1}{4m} \sum_{i,j} \left(A_{ij} - \frac{k_i k_j}{2m} \right) s_i s_j. \quad (3.7)$$

In the model investigated, vaccination occurs only through the first adoption of pro-vaccine opinion $N \rightarrow V_s$ (Sec. 3.2.1); as such, the number of changes of opinion Θ_* undertaken by agents

of either opinion can conceivably describe both the social and infection dynamics of the model and thereby yield warning signals of sudden transitions. Moreover, (3.2) shows that the probability of switching sentiment is sensitive to the number of infected agents in the individual neighbourhood; this dependence makes it highly likely that the probability of having an infected neighbour Γ_* is correlated not only with vaccine coverage both also with the rate of change of opinion throughout the length of each simulation.

3.2.3 Parameters and time series

Both (infection and social) layers of the network have an Erdős-Rényi random network structure

$$\mathcal{G}(10000, 0.003);$$

network size $D = 10000$ and mean degree $\langle d_n \rangle = 30$ were chosen for alignment with studies of similar coupled behaviour-infection models [274, 445], as well as for computational tractability. Each simulation starts with proportion $\alpha = 0.25$ of vaccinated pro-vaccine agents (states represented by pair (V_s, V_p)), with all others being susceptible anti-vaccine agents (pairs (N, S)). The probability of death per time step is μ , so that $\alpha \cdot \mu$ gives the probability of reset with initial state V_p and $(1 - \alpha) \cdot \mu$ the probability of reset with initial state S . $\xi = 1 \times 10^{-4}$ represents the probability of switching sentiment randomly. Noise parameter $\xi = 0.0001$ represents unsystematic fluctuations commonly seen in empirical studies [110, 301, 69]. The birth/death probability $\mu = 2.4 \times 10^{-4}$ affords an average life span of 80 years to each agent [97]. The case importation rate $\iota = 2.5 \times 10^{-5}$ of susceptible agents adds periodic impulses of infection as a test of resilience in non-endemic disease regimes. A complete list of variable notation and values is given in Tabs. 1.2 and 1.3.

Ensemble size was set at 20 for most parameter combinations used. Figure 3.2 shows nine time series from the infection and social dynamics for a succession of perceived vaccine risks κ ; all time series taken over all realisations showed an initial epidemic spread (defined in Sec. 1.2, examples in Figs. 3.2(a-c)), with subsequent decrease as the pool of susceptible agents is consumed. Noted is that Figs. 3.2d ($\kappa = 0.03125$) and 3.2e ($\kappa = 0$) both show almost perfect vaccine rates $\langle V_p \rangle$ over 49 realisations of the parameter values, while the corresponding social pro-vaccine consensus shown in Fig. 3.2g collapses as the perceived vaccine risk κ decreases into Fig. 3.2e. Finally, Figs. 3.2f and 3.2i ($\kappa = -0.03125$) show a median vaccination rate $[V_p] \approx 0.5$ with corresponding anti-vaccine consensus $[V_s] \approx 0$ respectively. Taken all together, these trends show that simple observation of the social dynamics may not allow for predictions of the infection dynamics and vice versa (as shown in Ch. 2). This is due to the permanence of the vaccine; an agent can change their vaccine opinion throughout their lifetime, but they cannot become ‘unvaccinated’ should their stances change.

The chosen parameter space $(\kappa, \sigma) \in [-1, 1] \times [0, 3]$ was sufficiently broad to capture transitions in both the infection (Fig. 3.3a) and social (Fig. 3.3b) dynamics. The contours in each panel of

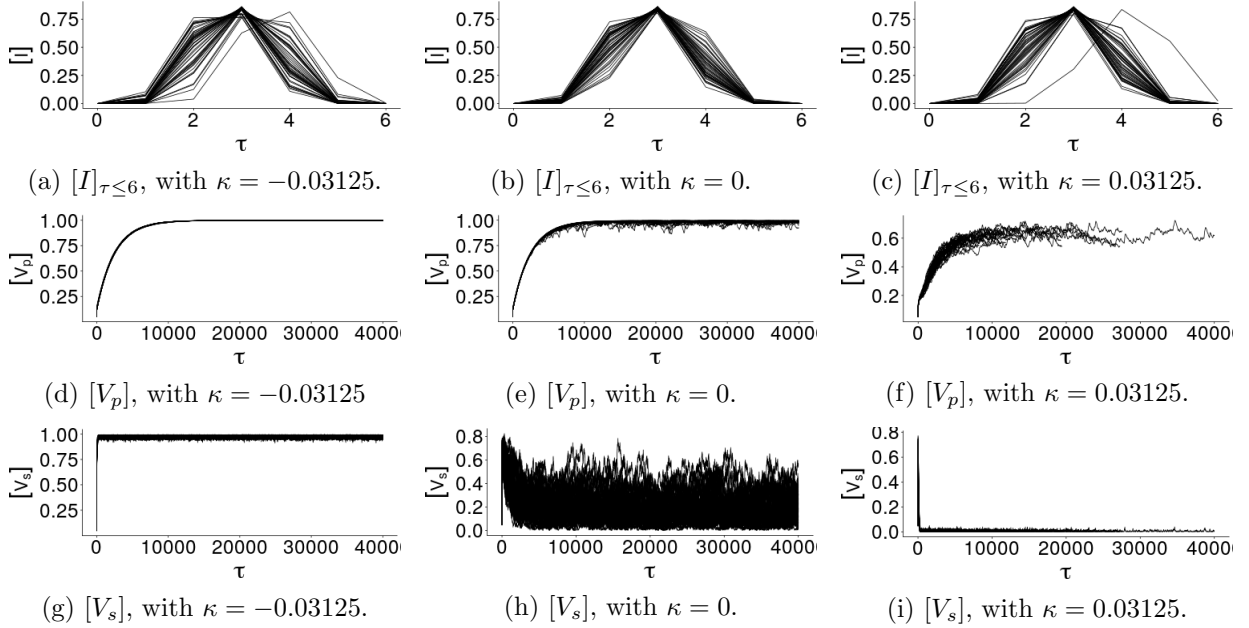


Figure 3.2: Example time series (49 realisations each) of social (d-f) and infection (g-i) dynamics of the model for social norm $\sigma = 0$, where time τ is measured in weeks (one time step is a week). Row (a-c) shows the initial epidemic spread of the disease over the first 6 time steps of each realisation. The changes in trend of $[V_s]$ with increasing perceived vaccination risk κ in (g,h) compared to the corresponding graphs of $[V_s]$ (d,e) show that the infection dynamics may not be a predictor of the social dynamics and *vice versa*.

Fig. 3.3 show the obvious correspondence between social (K_s) and infection (K_p) transitions and substantial changes in the clustering coefficient of the pro-vaccine subnetwork $\langle C_{V_s} \rangle$ (Fig. 3.3c) and the mean size of anti-vaccine communities $\langle |Z_N| \rangle$ (Fig. 3.3d).

3.3 Results

We again define the *infection transition* K_p as the value of perceived vaccine risk κ value at which the mean number of recovered agents in the model surpasses that of the vaccinated agents and vice versa, so that $\langle V_p \rangle \approx \langle R \rangle$; in Fig. 3.4, K_p is represented by the dashed vertical line marking the approximate intersection of the curves giving $\langle V_p \rangle$ and $\langle R \rangle$. Similarly, the *social transition* K_s is the κ value at which one of $\langle N \rangle$ (number of anti-vaccine agents) and $\langle V_s \rangle$ (number of pro-vaccine agents) surpasses the other so that $\langle V_s \rangle \approx \langle N \rangle$; this is marked by the dotted vertical line in Fig. 3.4 showing the approximate intersection of $\langle V_s \rangle$ and $\langle N \rangle$.




Figure 3.3: Contour plots of the region $(\sigma, \kappa) \in [0, 3] \times [-1, 1]$ of the parameter plane, capturing the transition dynamics of both the social and infection dynamics; σ represents the strength of injunctive social norms and κ the perceived health risk of the vaccine. The probability of infection is $p = 0.2$, with $\langle |Z_N| \rangle$ the mean size of anti-vaccine communities and $\langle C_{V_s} \rangle$ the clustering coefficient of the subnetwork of pro-vaccine agents.




Figure 3.4: Changes in trend in the κ -series of output variables $\langle R \rangle$, $\langle V_p \rangle$, $\langle N \rangle$, $\langle V_s \rangle$ predict both the social K_s (dotted vertical line) and infection K_p (dashed vertical line) transitions.

As in Ch. 2, the distance between the two transitions (i.e. $K_p - K_s$) is called the *intertransition distance*; as with similar models, we find that this intertransition distance decreases with increasing strength of the social norm σ . This is explicitly demonstrated by Fig 3.4, where an increase in the social norm from $\sigma = 0$ (Fig. 3.4, left) to $\sigma = 0.5$ (Fig. 3.4, right) brings the two vertical lines (representing transitions K_s and K_p respectively) together.





Figure 3.5: Graph showing the decreasing trend in the *intertransition distance* $K_p - K_s$. The inset graph shows the κ values at which the infection K_p (red) and K_s (blue) transitions occur. The decrease in this distance (irrespective of the individual values of K_s/K_p) potentially changes our interpretation of the EWS.


Figure 3.5 gives a full picture of the intertransition distance over the investigated parameter range $0 \leq \sigma \leq 2.2$; its decreasing trend with strengthening social norm σ is shown by the purple curve, with the inset panel showing the locations of the social K_s (blue) and infection K_p (red) transitions. The positivity of the graph tells us that $K_s < K_p$ for all strengths of the social norm $\sigma \leq 2.5$, so that the social transition always precedes the disease transition. Since physical vaccination is driven primarily by changes in sentiment, this is to be expected; strengthening social norms increases the alignment of behaviour and vaccine uptake, resulting in decreasing intertransition distance $K_p - K_s$. Warning signals can potentially indicate both transitions, or maybe only one of the two; this subtlety is lost with shrinking intertransition distance and so may impact the predictive power of any *early warning signal* tested.

3.3.1 Group size and census predict the social transition


Figure 3.6 shows the trends in the means of the numbers and sizes of communities and echo chambers (both pro- and anti-vaccine) at equilibrium. As for the number of echo chambers $\langle \#J_* \rangle$ (Fig. 3.6d), there is not much resemblance to the trend of the mean number of connected components $\langle \#Z_* \rangle$ (Fig. 3.6b); this is partly due to the low occurrence of echo chambers in this network. For $\sigma = 0$ (Fig 3.6d, left), there is no warning given by the mean number of pro-vaccine echo chambers $\langle \#J_V \rangle \approx 0$, whereas the mean number of anti-vaccine echo chambers $\langle \#J_N \rangle$ increases before K_s , reaching a maximum between the two transitions and reaching zero




(a) Mean community size; $\langle |Z_N| \rangle$ (anti-vaccine), $\langle |Z_V| \rangle$ (pro-vaccine).



(b) Mean number of communities; $\#Z_N$ (anti-vaccine), $\#Z_V$ (pro-vaccine).



(c) Mean echo chamber size; $\langle |J_N| \rangle$ (anti-vaccine), $\langle |J_V| \rangle$ (pro-vaccine).



(d) Mean number of echo chambers; $\langle \#J_N \rangle$ (anti-vaccine), $\langle \#J_V \rangle$ (pro-vaccine).

Figure 3.6: Trends of some measures of connectivity of the social network with respect to the perceived risk of vaccination κ . Vertical dashed and dotted lines representing the infection (K_p) and social (K_s) transitions (respectively) help to illustrate changes in the trends as transitions are approached. The strength of the social norm is $\sigma = 0$ for panels on the left and $\sigma = 0.5$ for panels on the right.

value approaching K_p . At first glance, the change in trend as the social norm strengthens to $\sigma = 0.5$ (Fig 3.6d, right) seems substantial.

Another view of the social dynamics involves analysing pairs of graphs. For example, Figs. 3.6a and 3.6b ($\sigma = 0$) together show the existence of a large community of pro-vaccine agents; for

$\kappa = -0.25$, there is a single community of pro-vaccine agents with size ≈ 9966.56 ,

$$\langle |Z_V| \rangle = 9966.56, \#Z_V = 1,$$

and a few **triads** of anti-vaccine agents ($\langle |Z_N| \rangle \approx 3.09$, $\#Z_N \approx 19.5$). This indicates the existence of one large pro-vaccine component interspersed with small triads of anti-vaccine agents. However, when perceived vaccine risk increases to $\kappa = 0.25$, there is a single anti-vaccine community of mean size 9992.8 ($\langle |Z_N| \rangle \approx 9992.83$, $\#Z_N \approx 1$) with a few individual pro-vaccine agents ($\langle |Z_V| \rangle = 1.82$, $\langle \#Z_V \rangle = 4.34$). These observations hold for the stronger social norm $\sigma = 0.5$ in Figs. 3.6a and 3.6b.


Intuitively, a similar pattern holds for the formation and erosion of echo chambers seen in Figs. 3.6c and 3.6d ($\sigma = 0$), and 3.6c and 3.6d ($\sigma = 0.5$). The scenario $\langle V_s \rangle \sim \langle N \rangle$ (for small κ) necessarily leads to the creation of a large community of pro-vaccine agents and thereby a large echo chamber; conversely, $\langle V_s \rangle \approx 1$ will result in small dispersed components of pro-vaccine agents with a small or empty interior, resulting in a very low number of echo chambers.

We use the Standard Normal Homogeneity Test (SNHT) from the `trend` package in R [454] to find a change point ($\text{SNHT}_\sigma\{\langle \Psi \rangle\}$) in Ψ , the κ -series of each EWS for each strength of the social norm σ ; the distance between the transition K_s and the change point is called the *lead distance*. Each panel of Fig. 3.7 gives the trends of lead distance


$$\text{Lead}_\sigma[\langle \Psi \rangle] = K_s - \text{SNHT}_\sigma\{\langle \Psi \rangle\} \quad (3.8)$$

of each EWS Ψ on the left, and a bar chart on the right showing the number of σ values for which each EWS gave the highest lead distance measured; positive trends indicate that the warning precedes the social transition. Each panel of Fig. 3.7 shows that most of the EWS shown give positive lead distances for $\sigma \leq 2.875$, and that most lead distances decrease as social norm σ strengthens.


Figures 3.7b and 3.7c expose all measurements of anti-vaccine echo chambers J_N as bad performers, along with the size of the smallest anti-vaccine community $\langle \min(|Z_N|) \rangle$ (Fig. 3.7a), the number of pro-vaccine communities $\langle \#Z_V \rangle$ (Fig. 3.7c) and the number of pro-vaccine echo chambers $\langle \#J_V \rangle$ (Fig. 3.7c); they all give very little warning of the social transition K_s for all strengths of the social norm σ , as well as giving the best warning (of all the EWS) for only $\approx 9\%$ of the tested range of the social norm σ . They all demonstrate significantly lower lead distances than those of the other EWS, giving the highest lead distance of all EWS for only 9% of σ values (as compared to $\geq 30\%$ for other EWS shown in Fig. 3.7). The performance of these EWS under different change point detection tests is shown in Figs. B.1-B.3. As social norm σ increases, pressure on each agent to conform to surrounding opinion becomes the main driver of self-organisation of the social dynamics. Since it also increases the speed of this transition between these opposing organised states, lead distances $\text{Lead}_\sigma[\langle \Psi \rangle]$ of EWS Ψ will decrease as the social norm σ increases.



(a) Communities: mean size of the largest $\langle \max(|Z_V|) \rangle$ and smallest $\langle \min(|Z_V|) \rangle$ pro-vaccine communities, and the largest $\langle \max(|Z_N|) \rangle$ smallest $\langle \min(|Z_N|) \rangle$ anti-vaccine communities, with the mean sizes $\langle |Z_N| \rangle$ and $\langle |Z_V| \rangle$.



(b) Echo chambers: mean size of the largest $\langle \max |J_V| \rangle$ and smallest $\langle \min |J_V| \rangle$ pro-vaccine echo chambers, and of the largest $\langle \max |J_N| \rangle$ and smallest $\langle \min |J_N| \rangle$ echo chambers, with the mean sizes $\langle |J_V| \rangle$ and $\langle |J_N| \rangle$.




(c) Mean numbers of groups: pro-vaccine $\#Z_V$ and anti-vaccine $\#Z_N$ communities, and pro-vaccine $\langle \#J_V \rangle$ and anti-vaccine $\langle \#J_N \rangle$ echo chambers.


Figure 3.7: Trends in the lead distance $\text{Lead}_\sigma[\Psi]$ for each EWS Ψ 's κ -series with respect to the strength of the social norm σ . The bar chart on the right of each panel shows the percentage of σ values for which the measured EWS gave the largest lead distance of all the EWS tested.

3.3.2 Clustering and uncertainty also provide early warnings


Figure 3.8 shows the trends in other EWS with respect to the social norm σ ; modularity of the social network $\langle Q_* \rangle$, the global clustering coefficient $\langle C_* \rangle$, number of sentiment changes $\langle \Theta_* \rangle$ and the probability of having an infected neighbour $\langle \Gamma_* \rangle$. Graph modularity $\langle Q_* \rangle$ and the number of sentiment changes $\langle \Theta_* \rangle$ give appreciable warning signals both transitions in Fig. 3.8, given their noticeable changes in trend approaching K_s (dotted vertical line) and K_p (dashed vertical line). However, the likelihood of having an infected neighbour $\langle \Gamma_* \rangle$ only predicts the disease transition when $\sigma = 0$; the probability of having an infected neighbour for both pro-vaccine $\langle \Gamma_V \rangle$ and anti-




(a) Network modularity. $\langle Q_N \rangle$ (anti-vaccine), $\langle Q_V \rangle$ (pro-vaccine), $\langle Q_{\Sigma} \rangle$ (total).




(b) Changes of opinion. $\langle \Theta_N \rangle$ (anti-vaccine), $\langle \Theta_V \rangle$ (pro-vaccine), $\langle \Theta_{\Sigma} \rangle$ (total).



(c) Probability of infected neighbours. $\langle \Gamma_N \rangle$ (anti-vaccine), $\langle \Gamma_V \rangle$ (pro-vaccine).



(d) Clustering coefficient. $\langle C_N \rangle$ (anti-vaccine), $\langle C_V \rangle$ (pro-vaccine), $\langle C_{\Sigma} \rangle$ (total).



(e) Clustering coefficient (no ribbon). $\langle C_N \rangle$ (anti-vaccine), $\langle C_V \rangle$ (pro-vaccine), $\langle C_{\Sigma} \rangle$ (total).

Figure 3.8: Trends of connectivity measures with respect to perceived vaccine risk κ . Vertical dashed and dotted lines represent the disease (K_p) and social (K_s) transitions (respectively). (e) shows the trends of the clustering coefficient without any indications of the standard deviation of the mean, for clarity. Social norm $\sigma = 0$ for the panels on the left, and $\sigma = 0.5$ for panels on the right.

vaccine $\langle \Gamma_N \rangle$ in Fig. 3.8c (for social norm $\sigma = 0$) only show substantial visible changes in trend approaching the disease transition K_p , unlike its uneventful approach to the social transition K_s .


This behaviour disappears when social norm $\sigma \rightarrow 0.5$ (Fig. 3.8c, right) due to the shrinking intertransition distance $K_p - K_s$; lead distances $K_p - \text{SNHT}_\sigma\{\langle \Gamma_* \rangle\}$ of both probabilities $\langle \Gamma_N \rangle$ and $\langle \Gamma_V \rangle$ for $\sigma = 0$ and $\sigma = 0.5$ are similar (0.95 and 0.98 respectively), showing that the distances between the warnings and the transitions K_p are similar in both cases. As σ increases to 0.5, K_s ‘moves closer’ to K_p ; this suggests that any warning of K_s occurs incidentally, rather than being directly caused by the model dynamics. This is intuitive; any measure of the probability of having an infected neighbour is an observation of the infection dynamics, therefore an assumption of some direct relationship between $\langle \Gamma_* \rangle$ and K_p is natural.

In Figs. 3.8d, the trends of the global clustering coefficient $\langle C_* \rangle$ are dwarfed by their standard deviations; Figs. 3.8e shows the trend without indicating the standard deviation of the mean, for clarity. In Figs. 3.8d and 3.8e, the respective clustering coefficients of the subnetworks formed by only anti-vaccine ($\langle C_N \rangle$) and pro-vaccine ($\langle C_V \rangle$) agents both show changes in trend approaching both K_s and K_p . The global clustering coefficient of the entire social network $\langle C_\Sigma \rangle$ holds constant value ≈ 0.03 , leading it to have the second-worst performance of all the EWS tested (Fig. 3.10).


Figure 3.9 shows the trends in lead distances of the remaining EWS. For all the EWS in Fig. 3.9, we can see that lead distance decreases with strengthening social norm σ , with all the EWS eventually failing (giving negative lead distances, so that warnings *follow* K_s - these are useless); modularity scores for the subnetworks formed by pro-vaccine $\langle Q_V \rangle$ (Fig. 3.9a) and anti-vaccine $\langle Q_N \rangle$ (Fig. 3.9a) agents give useful warnings for all social norms $\sigma \leq 2.40625$, while both the number of opinion changes $\langle \Theta_* \rangle$ (Fig. 3.9b) and the probability of having an infected neighbour $\text{Lead}_\sigma[\langle \Gamma_* \rangle]$ give useful signals for $\sigma \leq 2.90625$. As stated, the global clustering coefficient of the entire social network $\langle C_\Sigma \rangle$ (Fig. 3.9d) was undefined for most σ and negative for quite a few others, resulting in the worst performance of all the EWS tested and giving the highest lead distance for only 4% of the total range of the social norm σ . The subnetworks generated by pro-vaccine agents gave often unsubstantial though positive leads $\text{Lead}_\sigma[\langle C_V \rangle]$ (Fig. 3.9d), while the subnetworks formed by anti-vaccine agents (Fig. 3.9a) gave very little lead distance over most of the range of σ . The performances of these EWS under different change point detection tests are shown in Figs. B.5-B.7.

3.3.3 Finding the best and the worst EWS


We compare the EWS by finding the proportion of σ values for which each EWS gives the largest warning; this is shown in Fig. 3.10. We specify a *good* warning as one that gives the *highest* lead distance of all warnings for a single σ value and a *bad* warning as one that gives either a *minimal*, *negative* or *undefined* lead distance. For many σ values, the largest and smallest lead distances were not unique, so that the ratios on neither side sum to 100%.




(a) Network modularity score of pro- and anti-vaccine sub-networks.



(b) Number of sentiment changes experienced by pro- and anti-vaccine agents.



(c) Probability of a pro- and anti-vaccine agents having an infected neighbour.



(d) Global clustering coefficient of pro- and anti-vaccine sub-networks.

Figure 3.9: Trends in the lead distances of the remaining EWS with respect to the strength of the social norm σ . Compared to the other metrics shown here, the global clustering coefficient $\text{Lead}_\sigma[\langle C_* \rangle]$ performs badly.

Comparison of the panels of Fig. 3.10 gives a notion of ‘dependability’; were a restricted set of EWS to be employed, we would prefer *good* EWS (ones that give the largest lead distances) that




Figure 3.10: The proportions of social norm σ values for which each EWS gives the greatest (green) and least (red) warnings (biggest and smallest lead distances), corresponding to the charts shown in each panel of Figs. 3.7 and 3.9.

aren't also *bad* (giving the smallest lead distances). Some of the foremost EWS best satisfying this criterion are the counts of all different opinion changes $\langle \Theta_* \rangle$ and the mean size of anti-vaccine communities $\langle |Z_N| \rangle$ (30% best, 15% worst). The best performers are the modularity scores of the opinion subnetworks $\langle Q_V \rangle$, being the best EWS for 32% of σ values, with 18% bad warnings. Conversely, Fig. 3.10 suggests that the numbers of echo chambers $\langle \#J_* \rangle$ both bring up the rear, providing good warnings for only 1% of σ values while giving bad warnings for 98% of social norm σ values tested.

Overall, all observations of the sizes of *anti-vaccine* echo chambers $|J_N|$ on the network are poor EWS. Conversely, the sizes of *pro-vaccine* echo chambers $|J_V|$ perform generally well ($\geq 26\%$ best, 15% worst). Also, the mean and maximum sizes of anti-vaccine communities ($\langle |Z_N| \rangle$ and $\langle \max(|Z_N|) \rangle$ respectively) give two of the highest ratios of best warnings (both 30%) and not many bad warnings ($\leq 15\%$), but observations of the minimum size $\langle \min(|Z_N|) \rangle$ give only 1% good warnings and 52% bad warnings. The global clustering coefficient $\langle C_* \rangle$ performed particularly badly as an EWS, with $\leq 4\%$ best and 35% – 80% worst warnings.

Another observation made from Fig. 3.10 would be the relationship between the percentages of good and bad warnings; they are actually strongly anti-correlated, with a coefficient of -0.77 . The previously shown behaviours of the lead distances eliminate the possibility of an EWS that densely alternates between good and bad warnings, but some EWS do neither; for example, the total modularity of the network $\langle Q_\Sigma \rangle$ is trivially neither good (0%) nor bad (0%) by virtue of being everywhere undefined. Nontrivially, the clustering coefficient of the anti-vaccine subnetwork $\langle C_N \rangle$ gives intermediate warnings (neither good nor bad) for 60% of social norm strengths σ . Therefore, no choice need be made between minimising the percentage of bad warnings and maximising the percentage of good warnings; both strategies yield largely identical results.

Figure 3.11 shows the performance of each EWS per social norm value σ . Red tiles show where the EWS gave the smallest positive lead distance of all its peers, while green tiles represent the σ values for which the EWS gave the largest lead distance. Yellow columns show where all EWS gave equal lead distances and black tiles represent failed warnings (either undefined or negative lead distances). Therefore, the relative length of an EWS' red bar in Fig. 3.10 represents the percentage of that EWS' red tiles in its row in Fig. 3.11; the same correspondence holds between the length of an EWS' green bar in Fig. 3.10 and the percentage of green tiles in the related row in Fig. 3.11.

The main insight provided by Fig. 3.11 deals with patterns of performance; for instance, the overwhelming red colouration of the rows corresponding to the mean number of pro-vaccine communities and anti-vaccine echo chambers ($\langle \#Z_V \rangle$, $\langle \#J_N \rangle$ respectively) and all EWS related to anti-vaccine echo chambers J_N show that these EWS are quantitatively the worst of the group. Also, there is no detectable pattern in performance visible on the grid; in other words, the effectiveness of the EWS cannot be broken down by ranges of σ value. For higher values of the social norm $\sigma \geq 2.625$, the prevalence of yellow columns provides the observation that




Figure 3.11: Grid showing the relative performance of each EWS. Green tiles denote the σ values for which the EWS gave the highest lead distance, red tiles represent the smallest lead distance, grey tiles represent lead distances that are neither maxima nor minima and yellow tiles show where all EWS gave the same lead distances. Black squares represent failed warnings (negative lead distances) and white tiles represent undefined values (no lead distance).

performance seems not to vary as much among the EWS as it does for smaller σ values, but nothing else is immediately apparent.

For social norms $\sigma = 2.90625$ and $\sigma = 3$, none of the EWS give valid warnings; lead distances are all negative, except for the total clustering coefficient $\langle C_\Sigma \rangle$, the number of opinion changes by anti-vaccine agents $\langle \Theta_N \rangle$ and the network's total modularity score $\langle Q_\Sigma \rangle$ which are undefined. This confirms behaviour seen for large σ in Figs. 3.7 and 3.9; indeed, $\langle C_\Sigma \rangle$ is undefined for most social norms σ because of the disconnection in the social network. Total modularity $\langle Q_\Sigma \rangle$ is everywhere undefined (row of white tiles) for this reason.

3.4 Discussion

In this paper, we tested the use and effectiveness of different network measures as early warning signals (EWS) of sudden transitions in the social and infection dynamics of a multiplex model of

disease. For the parameter values used, we found that observations of the mean and maximum sizes of anti-vaccine communities appear to be the most effective EWS of all tested, unlike the size of the smallest anti-vaccine community (though it does give warnings signals); trends in the global clustering coefficient of the subnetworks formed by pro- and anti-vaccine agents respectively also marked these events, as well as the number of both respective communities and echo chambers preceding both transitions. This reflects the breakup of connected components on a network preceding a critical transition, an observation well supported by literature on percolation thresholds in random graphs.

A phenomenon of particular interest in this study was the formation and breakup of pro- and anti-vaccine echo chambers; we found that all observations of the sizes of *pro*-vaccine echo chambers (maximum, minimum and mean) performed well as warning signals, while observations of the sizes of *anti*-vaccine echo chambers performed poorly compared to other EWS. The modularity measure of the social network and the rate of opinion changes also warn of transitions of the social and disease transitions, representing changes in aggregate vaccine opinion and vaccine uptake crises respectively. As a direct observation of the infection dynamics of the model, the probability of having an infected neighbour (for both pro- and anti-vaccine agents) performed well as an EWS of vaccine crisis.

Through our proposal and study of effective graph connectivity measures, this study complements others in the field of early warning signals. A potential limitation to the study is our strict definition of an echo chamber; it remains to be seen whether different descriptions such as those featured in other studies of social media networks will result in a more effective or dependable EWS. Also, the graph connectivity measures seem suitable for tracking the dynamics of an evolving network; the inclusion of preferential link formation in the dynamics, as well as social and ‘on the ground’ interventions for different strengths of the social norm, present other interesting avenues of research. Finally, the inclusion of directionality of communication in the network may render the model more realistic [577, 309, 56, 382].

Together with other markers of spatial correlation and aggregation, the graph connectivity measures presented here contribute to the set of tools allowing us to leverage the ubiquity of social media involvement and the resulting data sets in the pursuit of adaptive strategies for maintaining public health.

Chapter 4

Spatial EWS are robust to delay and network sampling

Abstract

The ubiquity of hesitance sometimes defies social pressure, especially when individuals are required to make health-care decisions they deem momentous. In epidemiology, such intervention hesitance can both initiate and prolong infectious disease outbreaks, especially when paired with vaccine denial. Previous literature has yielded effective early warning signals (EWS) of disease outbreak and vaccine crisis for coupled behaviour-infection systems; these EWS arise from characteristic phenomena undergone by model dynamics during critical transition(s). In this study, we investigate the resilience of these EWS to the incorporation of a destructive delay to vaccination decisions. Simulations were conducted on a static small world network, using a model coupling an $SIRV_p$ infection model with a social dynamic resembling a voting game with abstention. We find that some of the EWS tested retain their efficacy despite fundamental differences in model behaviour. We also find that these EWS (both pair- and cluster-based) can be reliably used while observing as little as 60% of the total network with relatively small loss of accuracy. These findings not only show the resilience of these EWS, but also allow for potential expansion of use cases and reduction in computational resource requirements.

4.1 Introduction

When faced with important decisions, individuals combine values and the beliefs they hold as a guide to the action required [197]; common barriers to decision making are ‘information overload’,

unclear conception for the context of the decision, the importance of time constraints, perceived lack of the requisite knowledge and cognitive biases [197]. These barriers can lead individuals to ‘put off’ making a decision on an issue, which is disadvantageous in time-sensitive circumstances. Indecision around vaccination decisions is called *vaccine hesitancy* and is quite common, being reported by more than 90% of the world’s nations over the period 2015–2017 [318].

Vaccine hesitancy was defined by the Strategic Advisory Group of Experts on Immunization (SAGE) as “delay in acceptance or refusal of vaccination despite availability of vaccination services.” and states that “Vaccine hesitancy is complex and context specific, varying across time, place and vaccine. It is influenced by factors such as complacency, convenience and confidence.” [410] The SAGE working group also identified vaccine hesitancy as “being present when vaccine acceptance in a specific setting is lower than would be expected, given the availability of vaccination services.” The group then goes on to give a range of potential determinants for this hesitancy; the three categories are ‘contextual influences’, ‘individual and group influences’ and ‘vaccine/vaccination-specific issues’ [410].

In the context of this thesis, the major determinants of the dynamics of the models used so far (in Chs. 2 and 3) fall into this categorisation: the social norm (σ) is considered by the group as a *group influence* and the effect of the perceived adverse effect of the vaccine (κ) is considered a ‘vaccination-specific issue’. A more apt view interprets these two parameters as potentially encapsulating all of the reasons outlined by the working group in these two categories. The intentional simplicity of the (binary vote) models of social communication used in models V1, V2 and V3 precluded the treatment of more complex identified factors such as culture, religion and socio-economics. While useful as a guide, this view of hesitance is fundamentally phenomenological (not surprising, given the function and aims of the working group) and therefore is not directly applicable to our model; instead, we account for vaccine hesitance by adding a new decision category (*hesitance*) to the social dynamic used in Chs. 2 and 3, resulting in a ternary model. To our best knowledge, a similar evolving vaccine hesitance dynamic has not been previously added to an agent-based model in such an explicit way; one peripheral study features a preset level of vaccine hesitance in its model [93].

This account of vaccine hesitance exploits the framework of ternary voting games. Ternary voting games (TVGs) are generalisations of simpler voter models that add a third choice, referred to as ‘abstention’ [188, 429]; a *tripartition* is a map T from the set N of all agents to the decision set $\{-1, 0, +1\}$ (the truth values of the vaccination decisions available). In this case -1 represents anti-vaccine sentiment, 0 represents vaccine hesitance and $+1$ represents a pro-vaccine opinion. Our previous studies in Chs. 2 and 3 used the Erdős-Rényi (ER) random graph model $\mathcal{G}(N, p)$ [177]; this network structure is a common choice for opinion dynamic modelling [59, 524, 575, 110, 342], but has noted shortcomings [388]. Empirical studies have found that human social networks feature a high degree of clustering [5, 224], short path length and diameter [6], large connected components [304, 82], and highly skewed power-law degree distribution [183, 39, 305, 5, 79, 442] due to the existence of **hubs**, though these properties may change as the network evolves [297, 367].

While the ER model is not considered to have these features (even after tuning the probability p) [392], the attraction of the model is its exact solvability and the critical phenomena stemming from its construction [336, 343, 3].

These properties of high local clustering and small diameter are called *small world properties* and the networks that feature the are referred to as *small world networks* [367]; many studies show that these networks closely model social networks [4] such as those of scientific research collaboration [390], professional acting [18] and the world-wide web [5]. One construction frequently used in modelling opinion dynamics is the Watts-Strogatz model $WS(N, K, \beta)$, which is also referred to by some studies as ‘the small world network’; in that way, the phrase ‘small world’ refers both to graphs featuring small world properties (which can be random or scale-free) and specifically *Watts-Strogatz networks* themselves [388]. *Scale-free networks* feature a power-law degree distribution [367], are also used in social network simulation studies due to their perceived enhanced realism and similarity to empirical networks such as neural networks [79], power grids [442] and other networks [183, 39, 305, 5] (*preferential attachment* models can evolve, whereas the Watts-Strogatz construction is inherently static [388]). Scale-free networks do not generally have small world properties [1].

Data mined from online social networks is used in many ways, including notification of contemporary events [476, 444] and investigation of human behaviour and mood [339, 527, 556, 278, 368], a concept at the core of this thesis. However, the curation of social network data poses many problems. For one, the sheer size of current ever-growing social networks (such as Twitter) results in unfathomably huge data sets [450, 191, 552]. This makes sampling virtually unavoidable in social network analysis, and even then still poses huge computational challenges [450]. Many studies explore sampling technique for large social networks [450, 451, 27, 38], with the intention of representing the entire network somewhat closely [66, 94, 192, 539].

Here, we introduce vaccine hesitance to the social dynamics of the model **V1** through the addition of the *H* (hesitant) category/opinion, We also use a small world network of comparable size to those used in some previous studies for the purposes of comparison. Finally, we observe various randomly selected portions of the network and, via simulation, we gather *early warning signal* (EWS) from these portions. In Sec. 4.2, we concretely outline the changed social dynamics and give redefinitions of dynamics transition points and lead distance. Sec. 4.3 will present the results of simulation and lead distance analysis and finally Sec. 4.4 will summarise the relevance and context of the results, as well as outline potential limitations of the study.

4.2 Methods

As with Chs. 2 and 3, we assume the $SIRV_p$ natural disease history shown in Fig. 1.7a. However, our novel inclusion of vaccine hesitance to the previous dynamics gives us the NHV_s social dynamics of model **V4**, shown in Fig. 4.1.

4.2.1 Model V4

There are two major departures from the NV_s models (V1, V2 and V3) of Chs. 2 and 3: the first is the use of a small world network structure, and the second is that agents can no longer transition directly between the two extreme V_s (pro-vaccine) and N (anti-vaccine) opinions, but rather will first adopt a hesitant stance H before complete transition. These hesitant agents (H) can be influenced by both pro-vaccine and anti-vaccine social contacts while not themselves influencing the decisions of others, since their neutrality does not contribute to the effective strength of the surrounding social norm σ . Intuitively, hesitance does not inspire vaccination ($\not\rightarrow V_p$).

In each week, each agent n compares her opinion with some randomly chosen social contact a . Again, n can possibly reevaluate and change her opinion if agent a disagrees; the interaction of two extreme opinions may result in hesitance ($\rightarrow H$), while an hesitant agent will transition to a 's stance. For example, if a pro-vaccine (V_s) agent communicates with an anti-vaccine (N) agent and starts to doubt her current beliefs, she will become hesitant ($V_s \rightarrow H$); in this state, she is vulnerable to whomever she communicates with next. She could encounter the same (or other) anti-vaccine agent and be convinced to switch ($V_s \rightarrow H \rightarrow N$, over both weeks), or she could encounter a pro-vaccine agent and return to her previous supportive opinion ($V_s \rightarrow H \rightarrow V_s$, over both weeks). Alternately, she will not reevaluate her opinion if she communicates with an hesitant (H) agent, no matter her current opinion.

These new social dynamics are shown in Fig. 4.1. It can be said that the introduction of hesitance does not fundamentally change the structure (and function) of the infection model; from the point-of-view of infection spread, hesitant agents appear as anti-vaccine agents, in that neither of them vaccinate. Opinion changes inspired by interaction are not deterministic; we define the probabilities \mathbb{P}_n of opinion change for an agent n in the style of (1.1):

$$\mathbb{P}_n(a \rightarrow b) = \frac{1}{1 + \exp(-U_n^{a \rightarrow b})}, \quad (4.1)$$

where $a, b \in \{N, H, V_s\}$, $a \neq b$ are vaccine opinions. The indices $U_n^{a \rightarrow b}$ depend on the perceived risk of vaccine adverse effects (*perceived vaccine risk*) κ and the number of infected physical contacts I_n of agent n as:

$$U_n^{N \rightarrow H} = U_n^{H \rightarrow V_s} = -\sigma \left(\frac{d_n^N - d_n^{V_s}}{d_n} \right) - (\kappa - I_n), \quad (4.2)$$

$$U_n^{V_s \rightarrow H} = U_n^{H \rightarrow N} = -\sigma \left(\frac{d_n^{V_s} - d_n^N}{d_n} \right) + (\kappa - I_n). \quad (4.3)$$

The strength of an injunctive social norm is represented by σ (as in Chs. 2 and 3) and it acts to preserve the opinion currently held by the node; d_n and d_n^* respectively represent the total number of social contacts of agent n and the number of social contacts with the opinion $*$.




Figure 4.1: Representation of the opinion dynamics of model V4. In each time step, each agent may switch opinion from a to b with probability $\mathbb{P}_n(a \rightarrow b)$ upon interaction with a dissenting neighbour. ξ represents the noise parameter (that is, the probability of randomly switching opinion), while α is the probability of entering the simulation as a vaccinated pro-vaccine agent (V_s, V_p) and μ the probability of death/replacement.

The infection dynamics of this model are almost exactly those outlined in Sec. 1.3, except that agents may be assigned epidemiological state H with probability $1 - \alpha$, rather than state N as was done before (where α is the probability of entering the simulation as a pro-vaccine vaccinated agent). Therefore, Item 5 of the description given in Sec. 1.3 would here read

- ...
5. In each time step, every agent faces probability μ of death, at which they are instantly replaced by either a vaccinated pro-vaccine agent (joint state (V_s, V_p)) with probability α , or a susceptible hesitant agent (H, S) .
- ...

The flow of a single week is shown in Fig. 4.2, a flow diagram documenting the path of an agent n through the different phases of the model in each time step. Processes in yellow boxes occur only once throughout the **realisation** while processes in red boxes denote loops, with their instructions repeated for all agents. True/False tests and decisions are denoted by blue diamonds and grey boxes represent choices.

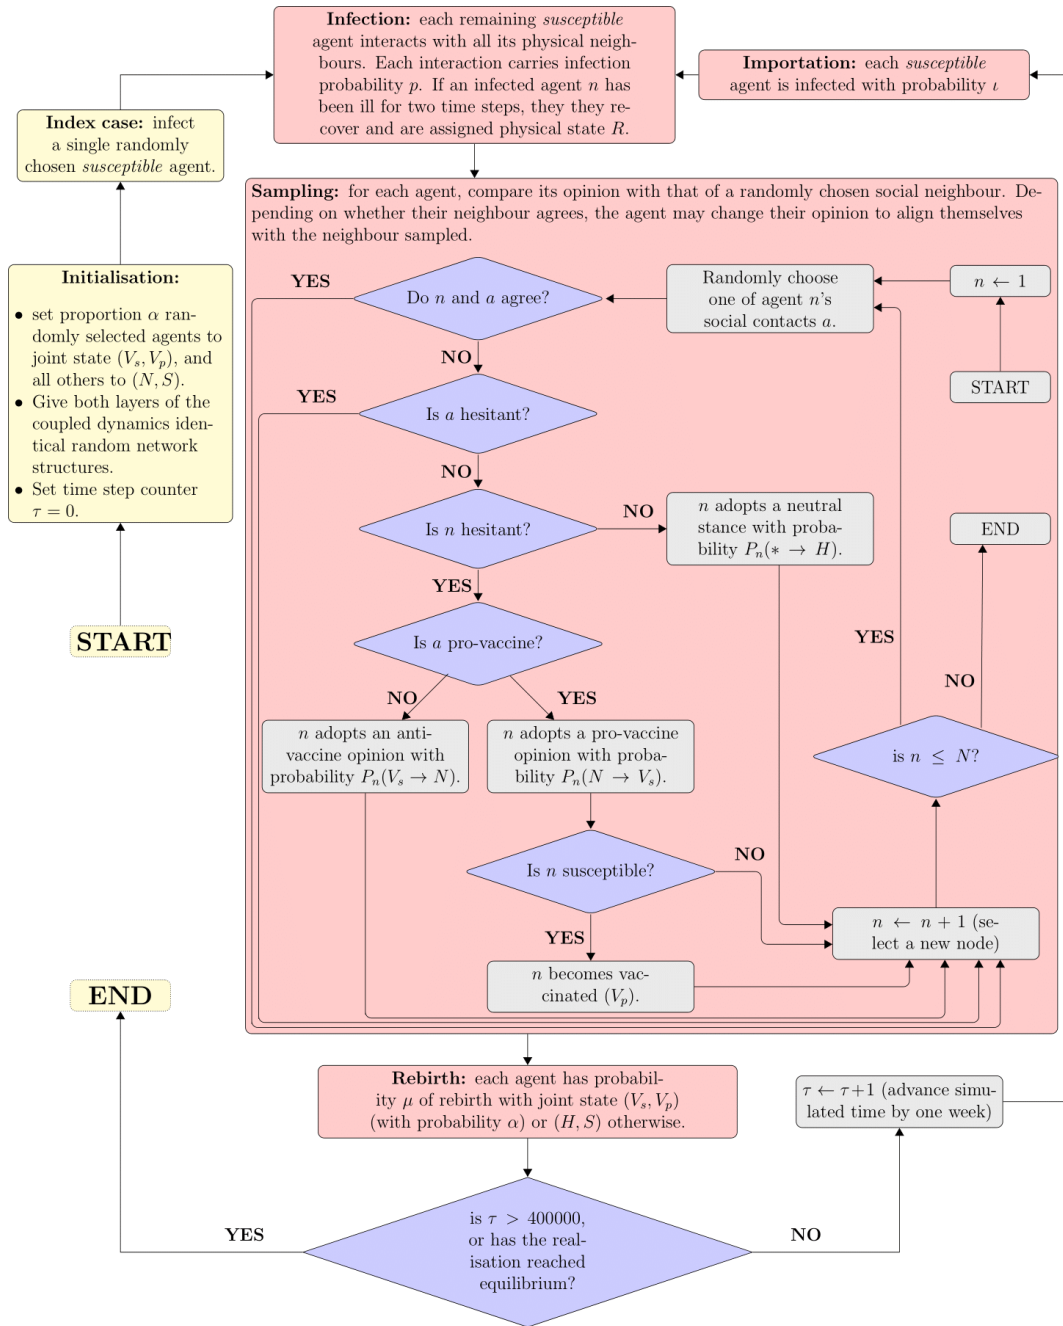


Figure 4.2: Implementation of the model dynamics used for each stochastic realisation.

4.2.2 Parametrisation

As with the previous studies in Chs. 2 and 3, the average life span of an agent in each simulation was set at 80 years, with $\mu = 2.5 \times 10^{-4}$. A small world network $\mathcal{WS}(p = 10^{-3}, D = 10000)$ was used to model agent connectivity in this study (in contrast to the previous random networks), with size $D = 10000$ and agents having $\langle d_n \rangle = 30$ neighbours each. A ratio $\iota = 1 \times 10^{-4}$ of susceptible cases were imported in each time step to provide periodic pulses of infection.

For the purposes of comparison, we retained the infection rate $p = 0.2$, with the probability of any agent randomly changing opinion kept at value $\xi = 1 \times 10^{-4}$. With this new model (V4), we found that the ranges $\sigma \in [0, 2.5]$ for the social norm and $\kappa \in [-1, 1]$ for the perceived vaccine risk allowed us to capture the transitions of the model, as shown in Fig. 4.3. As a further test of the efficacy and resilience of the EWS, we also compute each EWS by sampling proportion $\beta \in \{0.6, 0.8, 1\}$ of the (agents in the social) network.

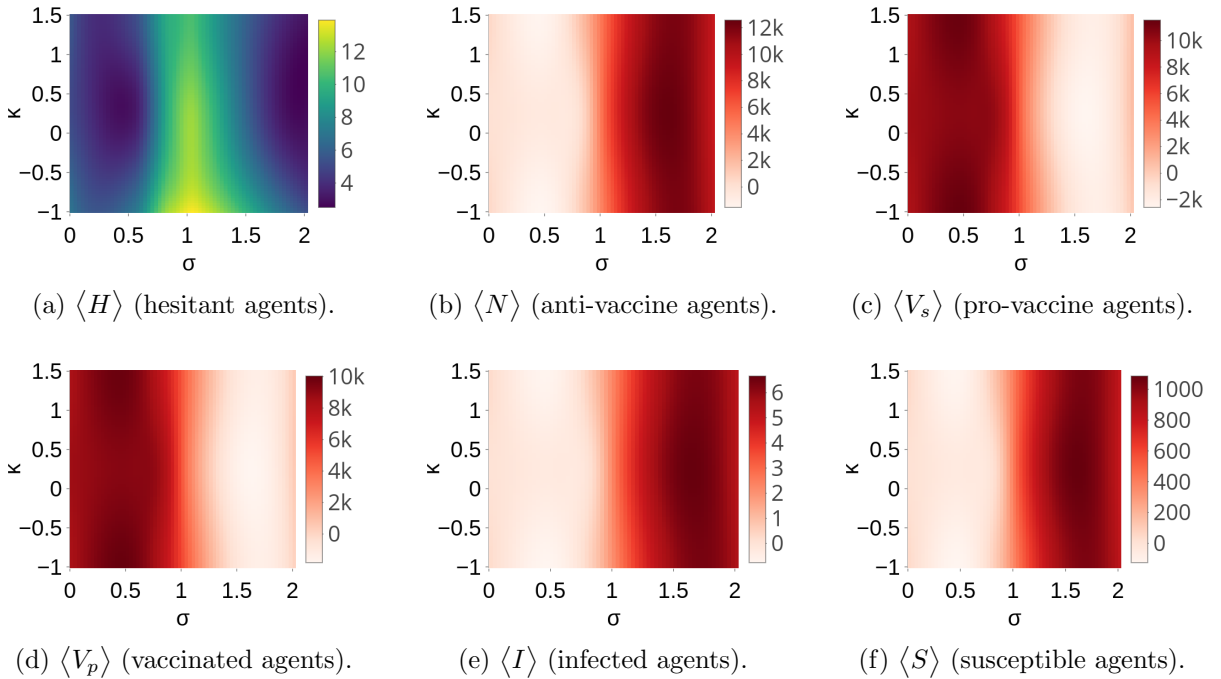


Figure 4.3: Heat map plots showing the ensemble means of model variables with respect to the social norm σ and the perceived vaccine risk κ . All agents were observed ($\beta = 1$).

An important feature of the model is its **bistability**, as demonstrated in Fig. 4.4; each point in each panel represents the mean value of the last 500 observations in the realisation, the criteria for which are outlined in Sec. 1.2. Fig. 4.4b justifies our choice of investigating a low initial

population of pro-vaccine agents ($\alpha = 0.25$), since we see that (generally) populations with a high initial proportion of pro-vaccine agents retain that proportion despite increasing risks of adverse vaccine effects. Our testing of EWS requires epidemic events and vaccine crises, so we judge higher initial proportions of pro-vaccine agents α to be inappropriate. However, our future discussion of clustering and [triadic closure](#) (Sec. 4.2.4) hints at α being an order parameter of the system controlling convergence in high risk (κ) consensus regimes, similar to a percolation threshold in our family of models.

This behaviour is also seen in model V3 (Fig. 4.4d) which also shows “stickiness” of the prevalence of vaccinated agents $\langle V_p \rangle$; for lower initial pro-vaccine populations ($\alpha \rightarrow 0$), intermediate perceived vaccine risks κ show model [bistability](#) before the establishment of either consensus regime $|\kappa| \rightarrow 1$, while higher initial pro-vaccine populations $\alpha \rightarrow 1$ survive this patch of instability intact. However, this behaviour is not seen for the number of susceptible agents $\langle S \rangle$ in either model V3 (Fig. 4.4d) or V4 (Fig. 4.4c), where the dynamics show instability both pre- and post-transition for all initial pro-vaccine population sizes αN .

4.2.3 Reconsideration of critical point estimates and lead distance


In Ch. 2.3, the transition point in the social dynamics K_s was defined as the first κ value at which the anti-vaccine and pro-vaccine curves crossed; that is, the κ for which $\langle N \rangle \sim \langle V_s \rangle$. Model V4’s inclusion of vaccine hesitance necessitates a change in this calculation; in Sec. 4.2, we remarked that hesitant agents appeared as anti-vaccine agents from the point-of-view of infection spread, since neither vaccinate. We reinforce this idea by defining the social transition (point) K_s^β as the κ value at which

$$\langle V_s^\beta \rangle \sim \langle N^\beta \rangle + \langle H^\beta \rangle, \quad (4.4)$$


where $\langle \Psi^\beta \rangle$ represents the mean proportion of agents with opinion/epidemiological state Ψ in a random sample of a proportion β of all agents. For example, $\langle V_s^{0.4} \rangle$ represents the proportion of pro-vaccine agents found in random sample of 40% of agents (in the social network), averaged over an ensemble of 20 independent realisations. For brevity, observations of the entire network will not be given a superscript, so that the true proportion of recovered agents on the network is written

$$\langle R \rangle \quad (= \langle R^1 \rangle).$$


Figure 4.5 shows the intertransition distance $K_p^1 - K_s^1$ for the model V4. Two stark differences from the trends in [intertransition distance](#) shown for models V1–3 (Fig. A.7) are that the distance actually increases with σ and is negative when $1.1 \leq \sigma \leq 1.15$. This shows that the infection outbreak actually preceded the vaccine crisis, but only by a small margin $|K_p^1 - K_s^1| \approx 0.0002$; however, given that these calculations used data resulting from multiple different realisations, we find it difficult to dismiss this phenomenon as a simple numerical artefact. These trends for other values of β are shown in Fig. C.1.




(a) Model V4, final number of recovered agents.



(b) Model V4, final number of vaccinated agents.



(c) Model V4, final number of susceptible agents.



(d) Model V3, final numbers of susceptible and vaccinated agents.

Figure 4.4: The total dynamics of models V3 and V4 shows bistability about the perceived vaccine risk $\kappa = 0$ for all parameter values tested.




Figure 4.5: Unsteady trend of the intertransition distance $K_p - K_s$ (purple), with the inset graph showing the estimated locations of K_s^1 (red) and K_p^1 (blue).

Intuitively, this is probably not completely due to the change in network structure (though that possibility cannot be completely dismissed), but mostly from our account of vaccine hesitance. By altering the social dynamics (not exerting social influence and delaying effective opinion change) while leaving infection dynamics virtually unchanged, hesitance changes the coupling of the social and infection dynamics. The link between the two dynamics revolves around the change to pro-vaccine stance (which results in vaccination), as shown by Fig. 4.2. Here, infection outbreak occurs not only through anti-vaccine consensus, but now also through widespread vaccine hesitance in the population.

(The opinion f) Vaccine hesitancy (H) was also cast in the role of an intermediate state between the two extreme pro- (V_s) and anti-vaccine (N) opinions and therefore acts to ‘slow down’ the social dynamics and delay the transition K_s^1 . Earlier graphs of the intertransition trend (Figs. 2.5a and 3.5) show that $K_p - K_s$ does not decrease smoothly; any consistent delay of K_s may result in briefly negative values, as well as much smaller intertransition distance in general. This is shown clearly by a comparison between Figs. 4.5 and 2.5a (model V2, no vaccine hesitance); the introduction of vaccine hesitance has compressed the median value of $K_p^1 - K_s^1$ by (roughly) a factor of 10. The trends in the number of pro- (V_s) and anti-vaccine (N) shown in both panels of Fig. 4.6 are superficially similar to those of previous models (Figs. 2.3a and 3.4), though the post- K_p^1 variation of the numbers of recovered (R) and vaccinated (V_p) agents is not familiar.

A final potential side effect of our introduction of hesitance to the social dynamics would be the complete disappearance of extreme (non-hesitant) opinions in consensus regimes, since it is now much more difficult to transition to an unpopular view. Since agents must now become hesitant and twice randomly choose a dissenting neighbour before completely switching opinion, any $N \rightarrow V_s$ (or $V_s \rightarrow N$) transition will take two weeks (two time steps, twice as long as in models V1-3), during the second of which the agent in question is much more likely to be drawn back to her previous position. One can well imagine that hesitant agents in a consensus regime




Figure 4.6: Trends of model variables of the social and infection dynamics for $\sigma = 0$ (left) and $\sigma = 0.25$ (right). Hesitance persists in the simulation across all strengths of the social norm σ .

in this model (V4) would have been identified instead as dissenting agents in previous models (V1-3). This may be manifested by the disappearance (undefinedness) of any trends sensitive to the total number of pro- or anti-vaccine agents in consensus regimes ($|\kappa| \approx 1$); an example is the probability of a pro-vaccine agent having an infected neighbour ($\langle \Gamma_V \rangle$), shown in Fig. 4.11b.




Figure 4.7: Estimated locations of the social and infection transitions K_s^β and K_p^β drift as the proportion of agents randomly sampled β changes.

We found that estimates of the transition points of both model dynamics obtained by sampling the agents of the network changed according to the size β of the sample, as shown in Fig. 4.7. Two perspectives can be defended here:

1. Estimates of K_s^1 and K_p^1 gathered from the entire network ($\beta = 1$) constitute a *ground truth* of the system, while the transition locations K_s^β and K_p^β gathered from subsets of the network are useful for comparison.
2. In a scenario where the entire network cannot be observed, the estimated transitions K_s^β and K_p^β are the only data possessed and so viewed as being the “real” transitions, though potentially specious.

This question is important to the calculation and interpretation of the *lead distance* $K - \Pi_*$; which is the true transition K in either dynamic? To avoid this ambiguity in this study, in cases

where the network is being sampled ($\beta < 1$), the transition locations will be taken as

$$K_s = \min(K_s^1, K_s^\beta), \quad K_p = \min(K_p^1, K_p^\beta), \quad (4.5)$$

by which we just use the earlier transition for lead distance calculations. We think that testing the ability of the EWS to predict the earlier transition in both dynamics is a better approach than making some consistent choice.

Finally, our main interest lies in predicting the first transition to occur independent of transition order. Since Fig. 4.5 shows that the social and infection transitions K_s^β and K_p^β do not occur in predictable order, lead distances are calculated here as the distance between the **change point** predicted by the **change point detection test** used and the earlier of the two transitions K_M , where

$$K_M = \min(K_s, K_p). \quad (4.6)$$

Here, we use the standard normal homogeneity test (SNHT) to get change points $\text{SNHT}_\sigma\{\langle\Psi\rangle\}$ for some EWS Ψ . We define a λ -network as the subnetwork (of the total social dynamics) formed by the set of agents with opinion λ . We then denote the lead distance of some EWS Ψ gathered from the λ -network formed by a random sample of proportion β of agents at social norm value σ as

$$\text{Lead}_\sigma[\Psi_\lambda^\beta] = K_M - \text{SNHT}\{\langle\Psi_\lambda^\beta\rangle\}. \quad (4.7)$$

For each value of the social norm σ , we will calculate which EWS give the largest lead distance and which give the smallest lead distance. Therefore, each EWS will have two attached percentages: the **LLT** denotes the proportion of values of the social norm σ for which the EWS gave the *largest lead distance* of all the EWS and the **ULT** will denote the proportion of social norm values σ for which the EWS in question gave *undesirable lead distances*, which refers to undefined lead distances, failures (negative lead distances) and minima (the smallest lead distances of all EWS). The effect of network sampling (*performance penalty*) on the lead distances of each EWS **Pen** will be quantified by taking the mean absolute (σ point-wise) difference between the lead distances of the EWS gathered from the entire network and from some simple random sample of proportion β of the network as a percentage of the mean lead distance. For example, the penalty of sampling proportion $\beta = 0.8$ of the network for the mutual information EWS $\langle\mathcal{M}\rangle$ will be written

$$\text{Pen}[\mathcal{M}^{0.8}] = 100 \cdot \frac{\langle|\text{Lead}_\sigma[\mathcal{M}] - \text{Lead}_\sigma[\mathcal{M}^{0.8}]\rangle}{\langle\text{Lead}_\sigma[\mathcal{M}]\rangle} \quad (4.8)$$

these results are also shown in Tabs. 4.1, 4.2, 4.3, 4.4 and 4.5.

4.2.4 Early warning signals: new and old

In this study, we reintroduce the following EWS:

- Sec. 2.2: Mutual information (\mathcal{M}), join counts ($[*, *]$),
 Sec. 3.2.2: Opinion communities (Z_*), echo chambers (J_*), global clustering coefficient (C_*),
 probability of having an infected neighbour (Γ_*), rate of opinion fluctuation (Θ_*),

and we omitted the following EWS:

- Sec. 2.2: Geary’s C (\mathcal{C}), Moran’s I (\mathcal{I}),
 Sec. 3.2.2: Network modularity (Q).

The above EWS candidates were excluded because Moran’s I was found to be a linear combination of join counts (App. A.1) and Geary’s C was found to be undependable (App. A.2.7 and Sec. 3.3 respectively). We include graph diameter and the number of (opinion-based) triads (described below) as candidate EWS.

Graph diameter is the length of the greatest geodesic between two vertices of a graph, otherwise stated as the maximum eccentricity of a vertex. Let \mathcal{G} be a graph with vertex set \mathcal{V} and a metric d . The eccentricity $\epsilon(u)$ of a vertex $u \in \mathcal{V}$ is the greatest distance between u and all other vertices

$$\epsilon(u) = \max_{a \in \mathcal{V}} d(u, a). \quad (4.9)$$

The diameter Ω_* is then the maximum eccentricity of the graph’s vertices, so that

$$\Omega_* = \max_{u \in \mathcal{V}} \epsilon(u) \left(= \max_{u, a \in \mathcal{V}} d(u, a) \right). \quad (4.10)$$

In this study, let Ψ be one of the vaccine opinions featured in the social (NHV_s) dynamics. Then we define the *diameter* of the Ψ community to be the greatest eccentricity of all agents with opinion Ψ :

$$\Omega_\Psi = \max_{u, a \in \Psi} d(u, a). \quad (4.11)$$

Triadic closure is a feature of social networks [580], and expresses the propensity of the formation of ‘triangles’ of agents; my friend’s friends are likely to become my friends as well, or alternately that two individuals are likely to be themselves friendly if they have a common friend [230]. The observation of this network feature, called *triad census*, counts the number of triads formed from potential triads; in the description of the calculation of the global clustering coefficient in Sec. 3.2.2, these were called T_Λ (the number of open triangles) and T_Δ (the number of closed triangles) respectively. In most studies, triad prevalence shows interconnection and homophily; here, the number of opinion triads instead speaks directly to the action of the social norm and nearest-neighbour influence.

In most (sociological) studies focused on the establishment of ties between actors, potential triads appear as open triangles; triads are eventually completed when the two unconnected individuals involved decide to acquaint themselves. This closure process however must be thought of differently because of the static nature of the network and our focus on individual opinion. Agents




Figure 4.8: Diagram showing the analogue of triadic closure occurring in model V4. Raj initially distrusts the MMR vaccine, but conversations with Ann and Ivo convince him of its benefit and safety, after which he adopts their opinion, thereby closing the triangle.

in model V4 are free to choose their own opinions but not their contacts (those are fixed); our problem is more akin to vertex colouring. A potential triad within this context is then a *dyad* - a neighbouring pair with the same vaccine opinion. Triadic closure occurs when the two agreeing nodes ‘convince’ the third agent to share their opinion as well, as shown in Fig. 4.8.

On the static social network shown in Fig. 4.8, the parents Ivo, Ann and Raj are Facebook friends that communicate regularly. Ivo and Ann are pro-vaccine (represented by a green hue) and actively encourage fellow parents to have their children vaccinated. Raj distrusts the MMR vaccine suggested by his family doctor and is hesitant (represented by grey hue). The social dynamics of model V4 specify that Raj doesn’t assert an opinion in his weekly chats with Ann and Ivo, but receives information from them (represented by directed edges reaching Raj), while Ann and Ivo exchange ideas civilly (joined by a solid edge). The function of the social norm σ (social pressure) acts on Raj during his chats with the gang and eventually he is convinced to have his daughter vaccinated (represented by green hue and solid edges joining the three since they now agree). This is the triad closure driven by the social norm; we denote the triad census for a vaccine opinion $*$ as Δ_* .

We remark then that the triad count is related to the similar join counts $[H, H]$, $[N, N]$ and $[V_s, V_s]$, though the relationship is not direct and not subject to the same interpretation. This is shown in Fig. 4.9, where panels (4.9a, 4.9b) together demonstrate an increase in the similar join count $[N, N]$ without a rise in the number of triads, and panels (4.9b, 4.9c) show a different number of triads with the same join count $[N, N] = 2$. While all three panels represent homophily, the long-term survival of a dissimilar opinion within social groups (such as H in Fig. 4.9a and V_s in Fig. 4.9b) suggests a weak or ineffective social norm σ .




Figure 4.9: Illustration of the tenuous relationship between the join count and the triad census.

4.3 Results

For model V4, the initial vaccine coverage $\alpha = 0.25$ ensured that all realisations showed epidemic spread (Sec 1.2) over the first 6 weeks (time steps) of each realisation, as shown in Fig. 4.10a-c.

Similar to the models of Chs. 2 (Figs. 2.2, A.2 and A.1) and 3 (Fig. 3.2), dynamics are highly sensitive to changes in perceived vaccine risk κ in the absence of social norm $\sigma = 0$. Fig. 4.10 shows that a steadily lessening perceived vaccine risk κ will bring the population from anti-vaccine consensus (Fig. 4.10d) to pro-vaccine consensus (Fig. 4.10f). Unlike previous results however (esp. in Ch. 2.3), vaccine uptake does not remain unchanged, as can be seen also in Fig. 4.10; $[V_p]$ increases with κ , growing from ≈ 0.45 (Fig. 4.10g) to ≈ 0.9 (Fig. 4.10i). This heavily foreshadows a change in the relationship between the social and infection transitions K_s and K_p , confirmed in Fig. 4.5.

4.3.1 Non-connectivity-based EWS retain their efficacy

In Fig. 4.11a, changes in the trend of the mutual information statistic $\langle \mathcal{M} \rangle$ predict the social K_s^1 (dotted vertical line) and infection K_p^1 (dashed vertical line) transitions for both values of the social norm $\sigma = 0$ (left) and $\sigma = 0.25$ (right). Similar to the previous models investigated (Figs. 2.4a, A.5 and A.6), $\langle \mathcal{M} \rangle$ increases to a maximum before K_s^1 and then falls to a post-transition minimum value.

The total number of opinion changes $\langle \Theta_\Sigma \rangle$ shows a similar trend (to $\langle \mathcal{M} \rangle$) around the transition(s) in Fig. 4.11c, with a fast increase to some pre- K_s^1 maximum value, decreasing through the transition to a post- K_p^1 minimum value. The probabilities of having an infected neighbour $\langle \Gamma_* \rangle$

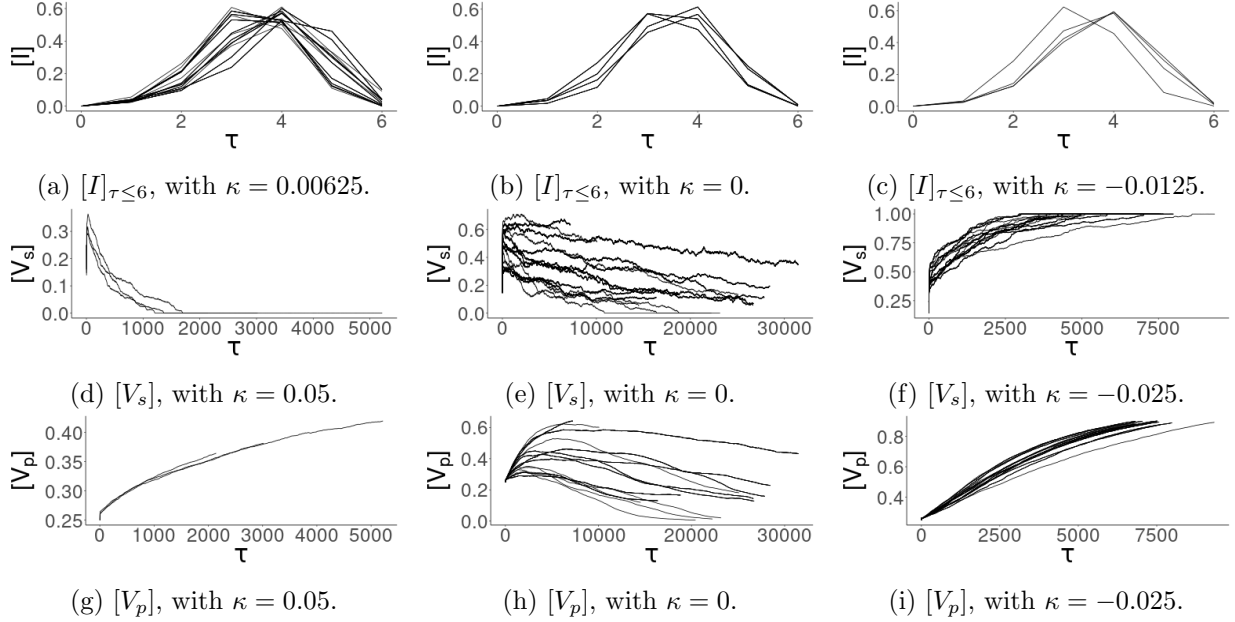


Figure 4.10: Time series demonstrating the sensitivity of system dynamics to small changes in perceived vaccine risk κ in the absence of a social norm ($\sigma = 0$). Time in week (time steps) is given by τ .

also predict the transitions, though the (changes in the) trends are not as clear as the previous EWS. For example, Fig. 4.11b shows that the probability of an anti-vaccine agent having an infected neighbour rises before the social transition K_s^1 for $\sigma = 0$ (left). However, the curve only begins shortly before $\kappa = -0.05$ (due to initial NaN values caused by the absence of anti-vaccine agents); similarly, the probability $\langle \Gamma_V \rangle$ is only periodically defined immediately after the transition K_p^1 . The probability of a hesitant agent having an infected neighbour $\langle \Gamma_H \rangle$ remains defined throughout the simulation due to the constant persistence of hesitant agents, shown in Fig. 4.6.

Figure 4.12 shows the lead distances of the probabilities of having an infected neighbour ($\langle \Gamma_* \rangle$, Fig. 4.12b) and the numbers of committed opinion changes ($\langle \Theta_* \rangle$ Fig. 4.12c) under the SNHT. Also shown in Fig. 4.12a are the lead distances of the mutual information statistic $\langle \mathcal{M} \rangle$ under the Lanzante, Pettitt, Buishand range and standard normal homogeneity tests.

Lead distances for the mutual information statistic are relatively high under the SNHT, ranging from 0.73 to 0.91 (Fig. 4.12a); sampling does not affect the lead distance much, since $0.52 < \text{Lead}_\sigma[\mathcal{M}^{0.8}] < 0.89$ (Fig. C.2b) and $0.57 < \text{Lead}_\sigma[\mathcal{M}^{0.6}] < 0.91$ (Fig. C.2a) represent small decreases with no marked change in the (admittedly jagged) appearance of the trend. The average lead distance of the mutual information dropped 11% when sampling 80% of the network ($\beta = 0.8$) and 9% when sampling 60% of the network ($\beta = 0.6$); see Tab. 4.1. An initial LLT of

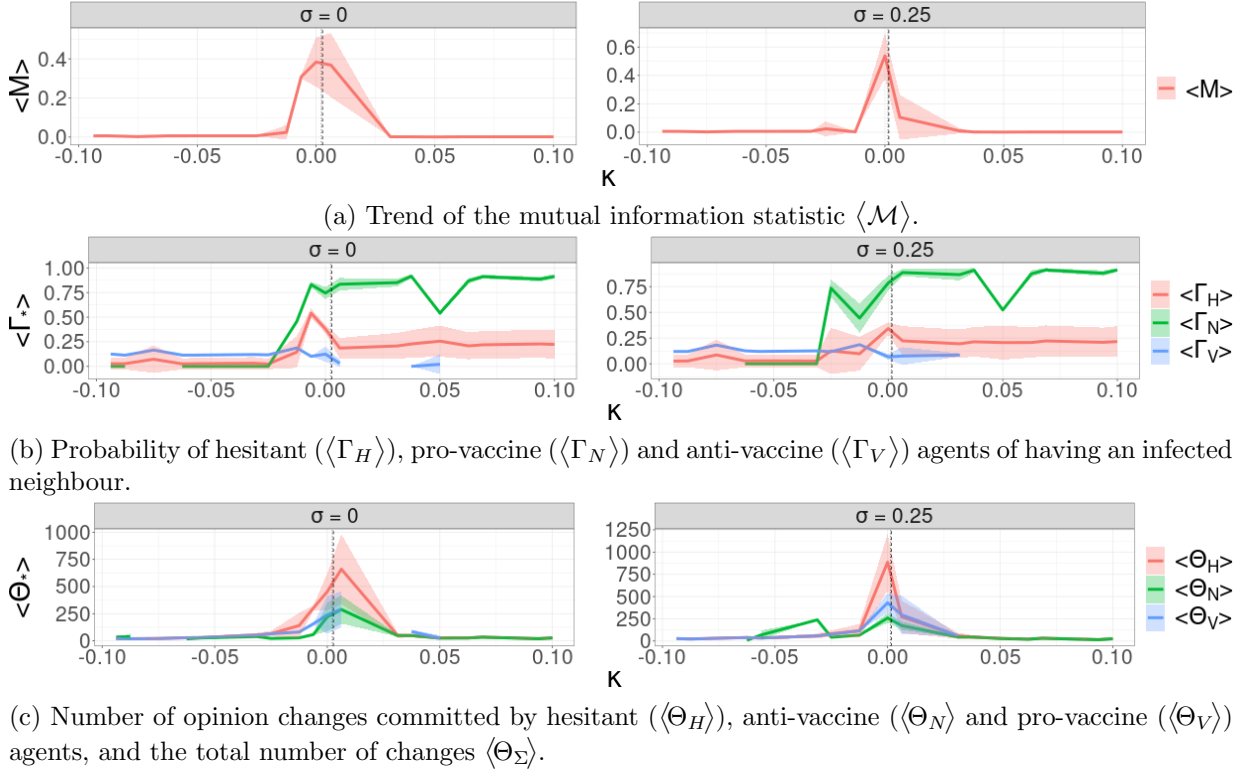



Figure 4.11: Changes in the trends of the κ -series of the mutual information statistic $\langle \mathcal{M} \rangle$, the probabilities of having an infected neighbour $\langle \Gamma_* \rangle$ and the numbers of opinion changes $\langle \Theta_* \rangle$ predict transitions in the infection (K_p^1 , dashed vertical line) and social (K_s^1 , dotted vertical black line) dynamics for both social norms $\sigma = 0$ (left panels) and $\sigma = 0.25$ (right panels).


33% (Fig. C.2c) dropped to 14% for $\beta = 0.6$ and 10% for $\beta = 0.8$.

The numbers of opinion changes $\langle \Theta_* \rangle$ also perform well under the SNHT (Fig. 4.12c); the total number $\langle \Theta_\Sigma \rangle$ and the number committed by pro-vaccine agents $\langle \Theta_V \rangle$ both have an LLT of 19%, with $\langle \Theta_H \rangle$ having an LLT of 10%. Again, lead distances don't change drastically when the network is sampled instead; $0.63 < \text{Lead}_\sigma[\Theta_*] < 0.98$, while $0.73 < \text{Lead}_\sigma[\Theta_*^{0.6}] < 0.92$ (Fig. C.2f) and $0.58 < \text{Lead}_\sigma[\Theta_*^{0.8}] < 0.83$ (Fig. C.2g); $\langle \Theta_H \rangle$ retains $\geq 10\%$ LLT for all sampling proportions β , while $\langle \Theta_\Sigma \rangle$ and $\langle \Theta_N \rangle$ both lose ground to other EWS (due to shifts in performance due to sampling). Penalties are similar, with the average lead distance declining 13% when sampling 80% of the network, and 9% when $\beta = 60\%$ of the network is sampled (Tab. 4.1).


Finally, the probability of having an infected neighbour $\langle \Gamma_* \rangle$ gave wildly varying lead distances (Fig. 4.12b), with a 5% LLT for both hesitant and pro-vaccine agents. In this way, there seemed to



(a) Lead distances of the mutual information statistic $\langle \mathcal{M} \rangle$ under the Lanzante, Pettitt, Buishand range and standard normal homogeneity tests.



(b) Probabilities of hesitant ($\langle \Gamma_H \rangle$), pro- ($\langle \Gamma_V \rangle$) and anti-vaccine ($\langle \Gamma_N \rangle$) agents having an infected neighbour.



(c) Opinion changes committed by pro-vaccine (V_s), anti-vaccine (N) and hesitant (H) agents.

Figure 4.12: Lead distances of non-connectivity-based EWS explored in the study, obtained by applying the SNHT. Bar charts on the right of the panels show the proportion of σ values for which each test gave the largest lead distance of all EWS tested.

be no immediately visible consistent change caused to the trend of $\text{Lead}_\sigma[\Gamma_*^\beta]$ as the sampling proportion β was varied (Figs. C.2d and C.2e). However, the average lead distance of $\langle \Gamma_V \rangle$ dropped $\geq 60\%$ upon sampling, with $\langle \Gamma_H \rangle$ suffering a similar fate (Tab. 4.1).

4.3.2 Group formation remains a strong EWS

Chapter 3 saw group structure and bonding on the network declared a strong family of EWS, and Figs. 4.13, 4.14, 4.15 and 4.18 show that the same holds in model V4 despite our changes to the social dynamic and the network structure.

Intuitively, the transition between pro-vaccine ($\kappa \approx -1$) and anti-vaccine ($\kappa \approx 1$) regimes causes the break-up of massive pro-vaccine communities (Fig. 4.13a) and the establishment of ever larger

EWS	β	Penalty (%)			Lead Distance (κ)		
		Minimum	Mean	Maximum	Minimum	Mean	Maximum
\mathcal{M}	0.4	0	12.63	40.51	0.39	0.76	0.89
	0.6	0	8.94	33.15	0.57	0.81	0.92
	0.8	0.49	11.27	36.34	0.58	0.75	0.86
	1	NA	NA	NA	0.73	0.85	0.92
$\langle\Theta_H\rangle$	0.4	0	11.97	34.79	0.7	0.83	0.95
	0.6	0	8.65	19.33	0.73	0.83	0.92
	0.8	0.51	9.82	30.41	0.58	0.75	0.86
	1	NA	NA	NA	0.64	0.81	0.92
$\langle\Theta_\Sigma\rangle$	0.4	0	8.96	22.57	0.7	0.82	0.95
	0.6	0	8.78	30.1	0.73	0.82	0.92
	0.8	0.5	12.95	44.64	0.58	0.74	0.86
	1	NA	NA	NA	0.7	0.83	0.98
$\langle\Theta_V\rangle$	0.4	0	8.96	22.57	0.7	0.82	0.95
	0.6	0	8.78	30.1	0.73	0.82	0.92
	0.8	0.5	12.77	44.64	0.58	0.74	0.86
	1	NA	NA	NA	0.7	0.83	0.98
$\langle\Gamma_H\rangle$	0.4	12.06	68.52	245.16	0	0.29	0.76
	0.6	0	71.48	241.14	0.01	0.3	0.79
	0.8	1.34	75.98	261.91	0.12	0.27	0.7
	1	NA	NA	NA	0.13	0.31	0.98
$\langle\Gamma_V\rangle$	0.4	13.19	70.29	153.01	0.01	0.52	0.95
	0.6	5.28	64.13	146.42	0.06	0.52	0.98
	0.8	0.88	61.08	171.92	0.14	0.49	0.86
	1	NA	NA	NA	0.19	0.47	0.95

Table 4.1: Lead distances and sampling penalties for mutual information $\langle\mathcal{M}\rangle$, the probability of having an infected neighbour $\langle\Gamma_*\rangle$ and opinion changes $\langle\Theta_*\rangle$ under the SNHT with respect to the sampling proportion β . A green-shaded row indicates that the EWS has mean sampling penalties less than 25% for all sample proportions.

anti-vaccine communities (Fig. 4.13b) as the perceived vaccine risk κ moves through the social and infection transition points K_s^1 and K_p^1 respectively. As was mentioned in Ch. 3, this formation and dissolution of community structure (and echo chambers, in Fig. 4.14) can be seen both as a cause and a necessary result of changing aggregate opinion for large $|\kappa|$, given that the underlying network is static; agents are not yet free to ‘move around’ in ways that resist dissolution. The only resistance in consensus regimes is through the preservation of individual opinion and, if an agent’s neighbour happens to achieve the same, then they will be counted as a group. The importance of connectivity measure lies in the region around the transitions K_* , where aggregate opinion doesn’t necessarily determine group settlement.

Small pockets of hesitant agents persist at most perceived vaccine risk values κ and experience quick growth as a transition K_* is approached. The rough symmetry of the trends in $|Z_H|$ (Fig. 4.13c) is to be expected from the intermediate role played by the stance H in the social

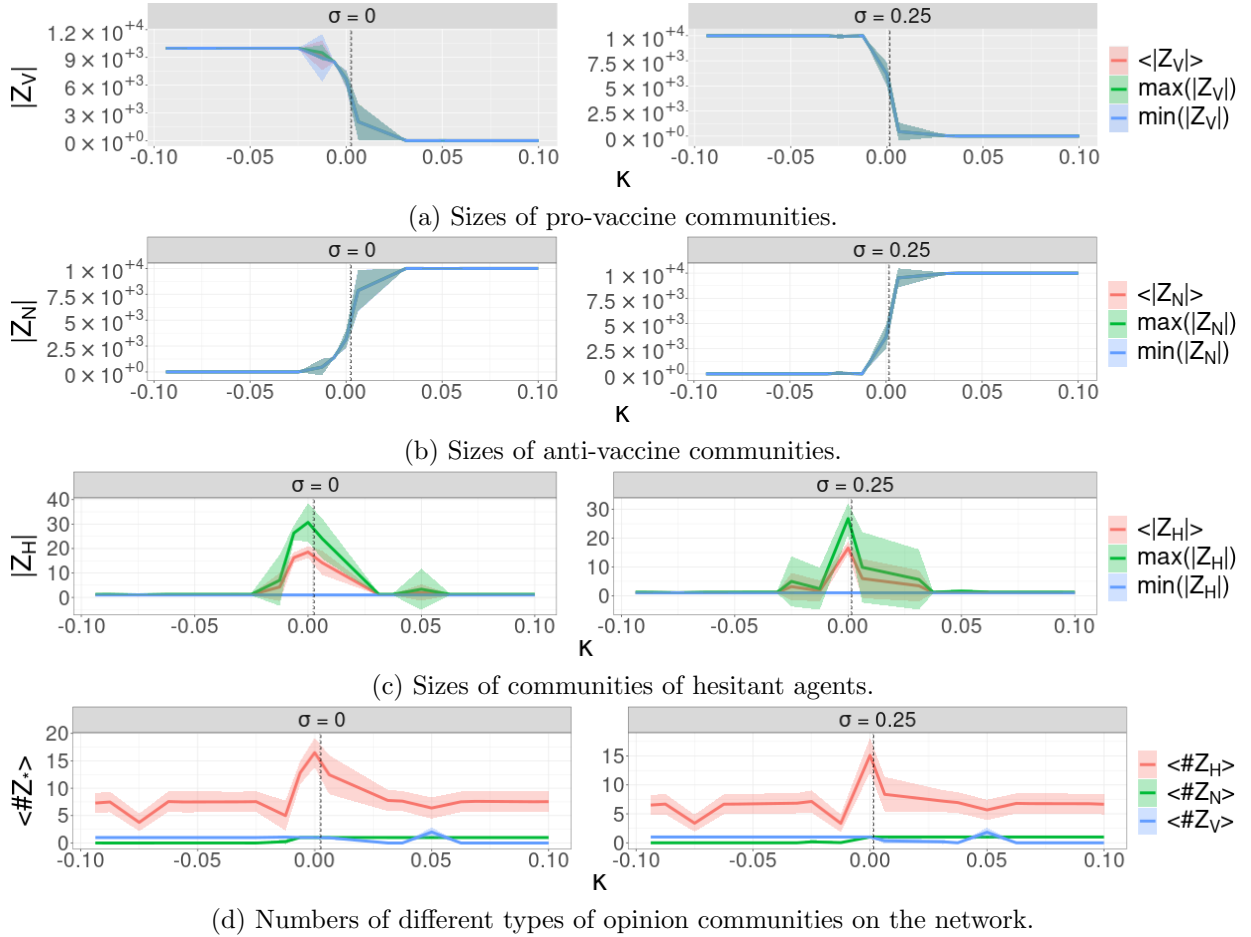


Figure 4.13: Sudden changes in trend of the sizes ($|Z_*|$) and numbers ($\#Z_*$) of opinion communities with respect to the social norm σ strongly signal upcoming transitions in the model dynamics. Social norm $\sigma = 0$ for panels on the left and $\sigma = 0.25$ for panels on the right.

dynamics. Figure 4.13d shows a consistently higher number of hesitant communities ($\#Z_H$) than other types, but the other panels of Fig. 4.13 show that the sizes of these communities are far smaller. This confirms the low level of persistence of opinion H shown in Fig. 4.6 for all values of perceived vaccine risk κ .

The trends of the sizes and numbers of echo chambers in Fig. 4.14 are broadly similar to those of the sizes and numbers of communities in Fig. 4.13, and so they also give clear warnings approaching the transition. The large sizes of pro-vaccine ($|J_V|$) and anti-vaccine ($|J_N|$) echo chambers in consensus regimes ($|\kappa| \approx 1$) represent the size of a single giant component made of

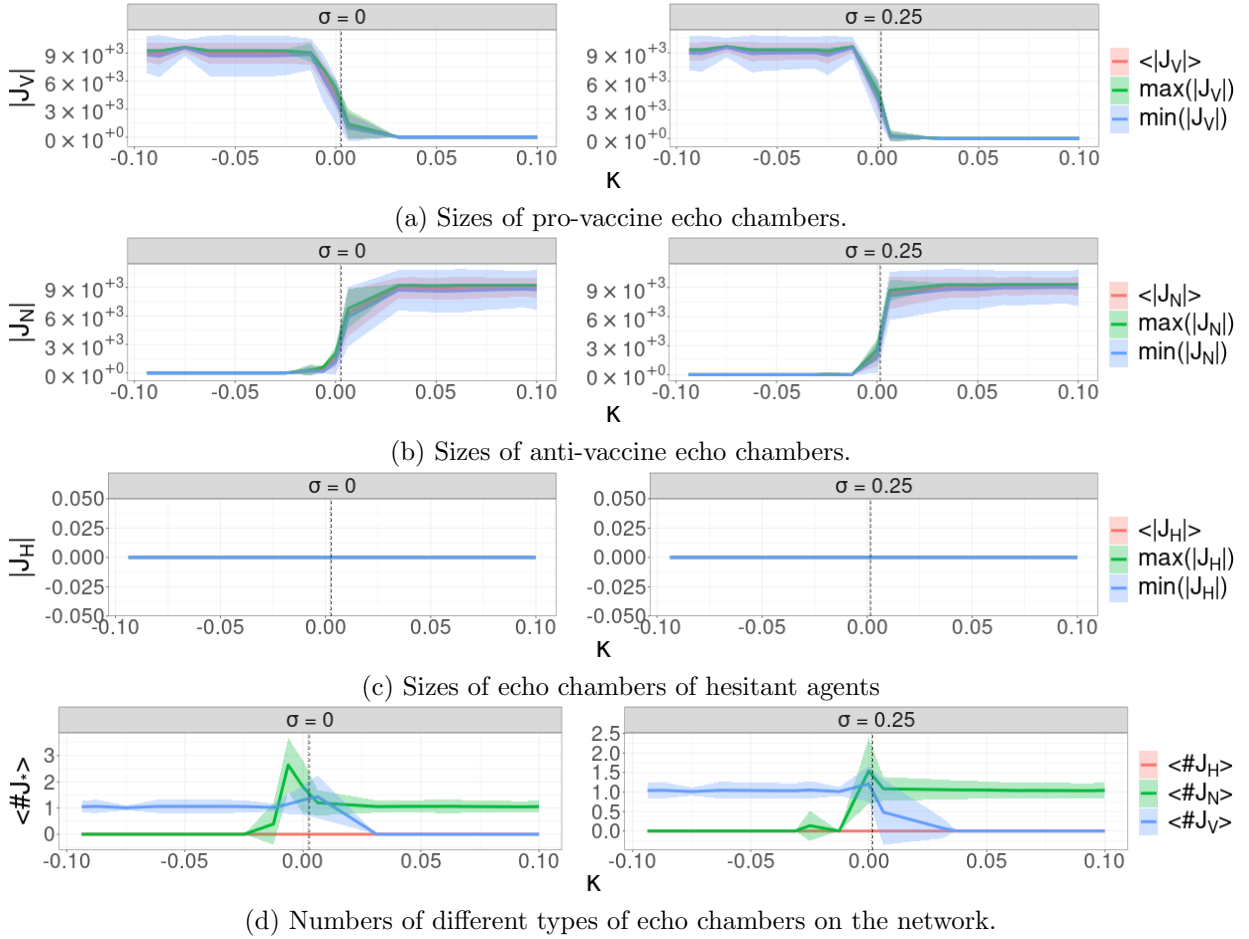


Figure 4.14: Changes in the sizes ($|J_*|$) and numbers ($\langle \#J_* \rangle$) of echo chambers in the network warn of the transitions K_s^1 and K_p^1 . Social norm $\sigma = 0$ for the panels on the left and $\sigma = 0.25$ for panels on the right.

interior agents of the large opinion component on the network. Figure 4.14d shows a low number of all types of echo chambers on the network; this is alternately due to the existence of a single large opinion community in a consensus regime, or scattered opinion close to a transition. Fig. 3.1 shows how moderately sized communities can possibly host only small echo chambers (if any at all), as a consequence of the restrictive definition given in Sec. 3.2.

One other observation is the universal absence of echo chambers of hesitant agents (Fig. 4.14c). As stated in the Methods (Sec. 4.2), hesitant agents do not exert any influence over their contacts (shown by the absence of a d_n^H term in (4.2)). Since this neighbour interaction is the major

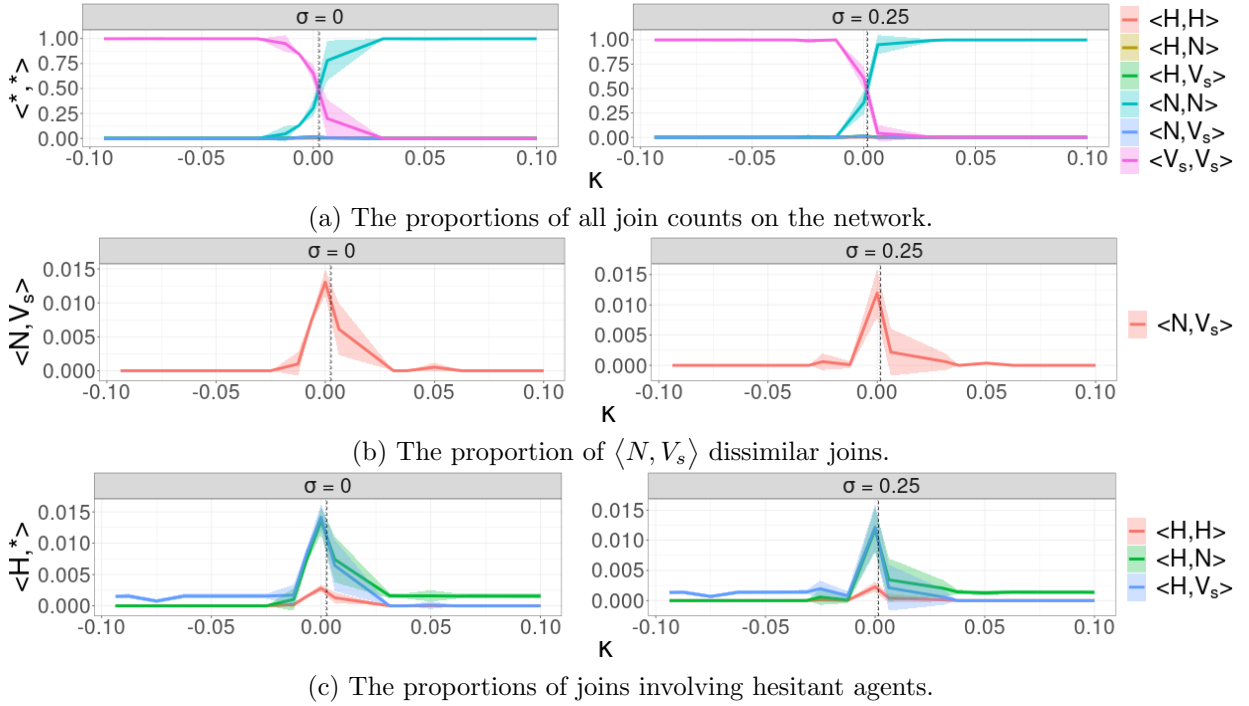
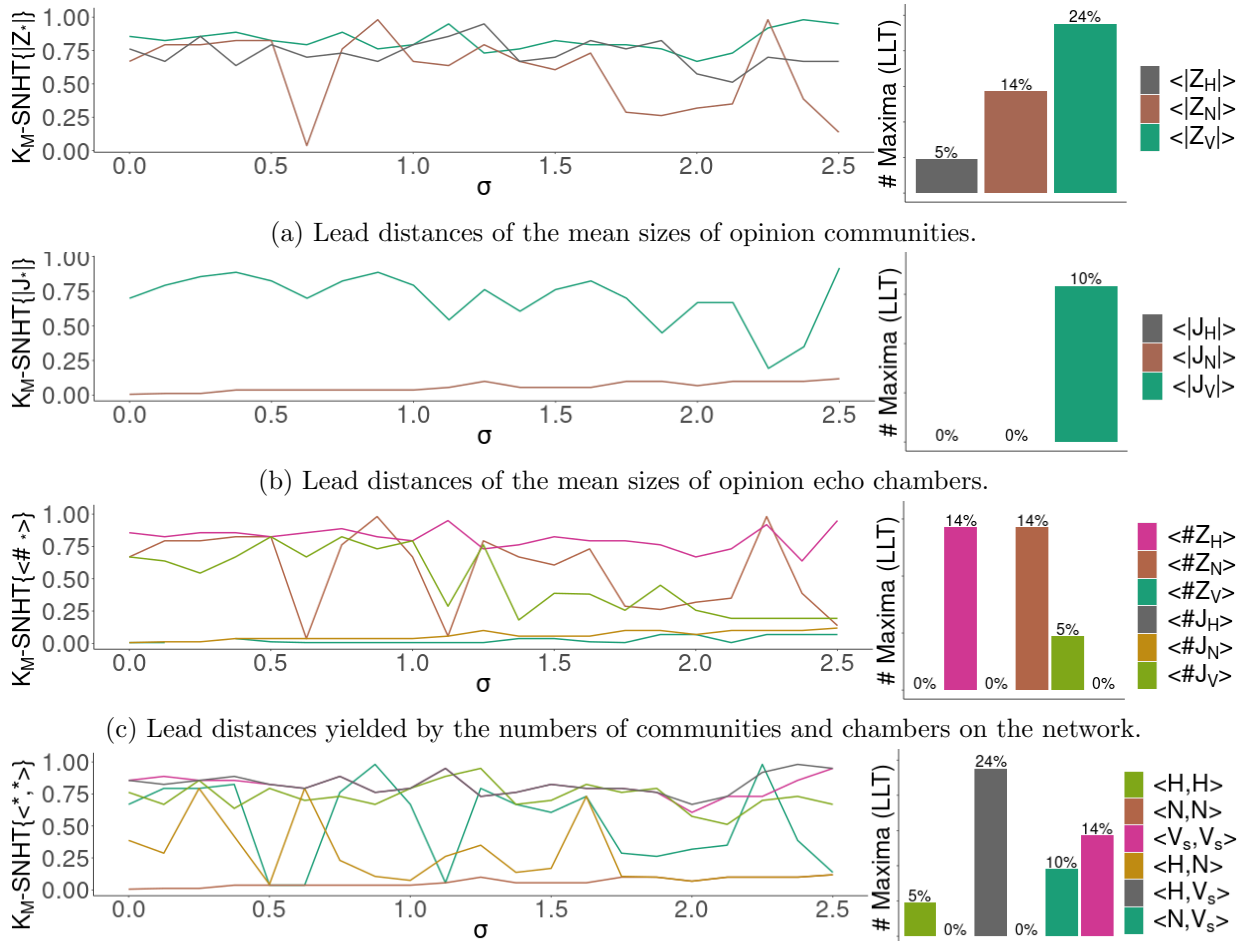


Figure 4.15: Trends in the joint count statistics of the network also give clear warnings of approaching social and infection transitions. Social norm $\sigma = 0$ for panels on the left and $\sigma = 0.25$ for panels on the right.

driver of grouping in consensus regimes, any collection of hesitant agents is quickly invaded and extremised without this strong imitation dynamic. Such echo chambers may appear briefly over the time evolution of the realisation, but do not persist to equilibrium.

Finally, Fig. 4.15 shows the trends of the joint count statistics of the social dynamics. Clearly, they all give clear warnings of the transitions K_* . As is expected, Fig. 4.15a shows that consensus regimes are dominated by similar joins between agents with the prevalent opinion. The pro/anti dissimilar join count $\langle N, V_s \rangle$ (Fig. 4.15b) still strongly signals the upcoming social and infection transitions via a roughly symmetric step pre-transition increase with the achievement of a maximum close to the transition. The proportion of joins involving hesitant agents $\langle H, * \rangle$ (Fig. 4.15c) also behaves similarly to $\langle N, V_s \rangle$. As with the previously mentioned sizes ($|Z_H|$, Fig. 4.13c) and numbers ($\#Z_H$, Fig. 4.13d) of hesitant communities, hesitancy's role in the social dynamics leads to the expected rising-falling trend roughly symmetric about the transitions K_* .

Figure 4.16 shows the trends in the lead distances yielded by some of the connectivity-based EWS under the SNHT, with respect to the social norm σ . The strongest performers of the mean



(d) Trends in the lead distances of the various join count statistics. Bar charts on the right give the percentage of social norm σ values for which the EWS gave the largest lead distance of all EWS.

Figure 4.16: Trends of the lead distances of connectivity-based EWS under the SNHT with respect to the social norm σ . Bar charts on the right of the panels show the proportion of σ values for which each test gave the largest lead distance of all EWS tested. Quite a few indicators perform poorly, with near-zero lead distances.

community sizes $|Z_*|$ are those of the anti-vaccine ($|Z_N|$) and pro-vaccine ($|Z_V|$) communities in Fig. 4.16a, with LLT 10% and 24% respectively. EWS related to chambers of hesitant agents ($|J_H|$ and $\langle \#J_H \rangle$) do not perform well, while the $\langle H, V_s \rangle$ dissimilar join count performs the best of all the join counts (Fig. 4.16d), with LLT 24%. The poorest-performing EWS were mostly connected to anti-vaccine sentiment; the sizes and number of anti-vaccine echo chambers $|J_N|$ and $\langle \#J_N \rangle$, and the anti-vaccine similar join count $\langle N, N \rangle$ all yielded consistently low lead distances and LLT under the SNHT.

Changing the size $\beta \cdot N$ of the random sample slightly alters the performance of the EWS here. For instance, the mean sizes of the various opinion communities $\langle |Z_*| \rangle$ all have LLT 10% when $\beta < 1$ (Figs. C.3a and C.3b). Though the total range of lead distances was not affected ($0 < \text{Lead}_\sigma[|Z_*|^\beta] < 1$ for all β), $\langle |Z_H| \rangle$ and $\langle |Z_N| \rangle$ saw greatly lowered lead distances when $\beta = 0.6$ (Fig. C.3a) and $\beta = 0.8$ (Fig. C.3b); for example, the mean lead distance $\text{Lead}_\sigma[|Z_N|]$ falls by 72% when only 60% of the network is sampled, and average lead distance $\text{Lead}_\sigma[|Z_H|]$ drops 30% when 80 of the network is used (Tab. 4.2).

EWS	β	Penalty (%)			Lead Distance (κ)		
		Minimum	Mean	Maximum	Minimum	Mean	Maximum
$\langle Z_H \rangle$	0.4	4.28	21.91	47.12	0.45	0.71	0.95
	0.6	0	22.93	55.69	0.29	0.6	0.95
	0.8	3.71	35.06	81.68	0.14	0.53	0.89
	1	NA	NA	NA	0.51	0.73	0.95
$ J_N $	0.4	0	42.08	114.35	0	0.03	0.09
	0.6	0	19.31	72.77	0.01	0.06	0.12
	0.8	6.91	63.54	97.04	-0.05	0.02	0.12
	1	NA	NA	NA	0.01	0.06	0.12
$\langle Z_N \rangle$	0.4	0	35.31	124.3	0.14	0.57	0.95
	0.6	0	57.94	138.61	0.04	0.51	0.98
	0.8	6.23	57.32	145.57	0	0.54	0.92
	1	NA	NA	NA	0.04	0.57	0.98
$\langle J_V \rangle$	0.4	0	20.22	65.87	0.46	0.74	0.95
	0.6	0	27.77	116.62	0.1	0.66	0.95
	0.8	0.59	28.29	111.57	0.14	0.64	0.92
	1	NA	NA	NA	0.19	0.7	0.92
$\langle Z_V \rangle$	0.4	0	13.47	52.2	0.45	0.74	0.92
	0.6	3.78	13.33	56.74	0.48	0.79	0.92
	0.8	0.5	10.9	41.1	0.55	0.77	0.92
	1	NA	NA	NA	0.67	0.83	0.98

Table 4.2: Lead distances and sampling penalties of the mean sizes of opinion communities $\langle |Z_*| \rangle$ and echo chambers $\langle |J_*| \rangle$ for different sampling proportions β , using the SNHT. A green-shaded row indicates that the EWS has mean sampling penalties less than 25% for all sample proportions.

The lead distance trends of the mean sizes of echo chambers $\text{Lead}_\sigma[|J_*|^\beta]$ were not greatly affected as β was changed, except that the lead distances $\text{Lead}_\sigma[|J_V|^{0.8}]$ saw greater fluctuation for larger

values of the social norm σ (in Fig. C.3d) than with other sampling proportions β (Figs. 4.16b and C.3c). All $\langle |J_V| \rangle$ are heavily affected by sampling, as seen in Tab. 4.2; for instance, the average lead distance given by the average size of pro-vaccine communities falls by 27% when only 60% of the network is sampled (i.e., $\text{Pen} [|J_V^{0.6}|] = 24.5$). $\text{Lead}_\sigma [|J_N|^{0.8}]$ is undefined for $\sigma < 1.25$ (Fig. C.3d), unlike its low value with other values of β (Figs. C.3c).

EWS	β	Penalty (%)			Lead Distance (κ)		
		Minimum	Mean	Maximum	Minimum	Mean	Maximum
$\langle \#Z_N \rangle$	0.60	0.00	70.33	159.86	0.04	0.49	0.98
	0.80	0.94	60.56	175.72	-0.00	0.47	0.92
	1.00				0.04	0.60	0.98
$\langle \#Z_V \rangle$	0.60	0.00	114.02	299.30	-0.01	0.02	0.06
	0.80	-19.76	-229.29	-652.12	-0.08	-0.02	0.04
	1.00				-0.01	0.03	0.07
$\langle \#Z_H \rangle$	0.60	0.00	12.37	58.68	0.42	0.75	0.95
	0.80	0.61	15.25	46.54	0.49	0.73	0.86
	1.00				0.64	0.81	0.95
$\langle \#J_N \rangle$	0.60	0.00	22.61	77.30	0.01	0.06	0.12
	0.80	46.07	241.03	330.88	-0.05	0.02	0.06
	1.00				0.01	0.06	0.12
$\langle \#J_V \rangle$	0.60	5.24	38.41	100.69	0.06	0.60	0.95
	0.80	1.01	54.04	144.90	0.07	0.44	0.92
	1.00				0.19	0.51	0.82

Table 4.3: Lead distances and sampling penalties of the numbers of opinion communities $\langle \#Z_* \rangle$ and echo chambers $\langle \#J_* \rangle$ for different sampling proportions β , using the SNHT. A green-shaded row indicates that the EWS has mean sampling penalties less than 25% for all sample proportions.

The lead distances of the numbers of opinion communities and echo chambers maintain their wide range for all values of β (Figs. C.3e and C.3f), with the number of communities of hesitant agents $\#Z_H$ maintaining high median lead distance for both sampling proportions β . Of this group, the number of hesitant communities sustained the lowest sampling penalties, as seen in Tab. 4.3; the smallest dip in performance occurred for the number of H communities ($\text{Pen} [\#J_H^{0.8}] = 15.25\%$), occurring when 80% of the network was sampled.

The lead distances of the join counts show less variation for smaller values of the social norm σ when the network is sampled (Fig. C.3g and C.3h) than for the ‘true’ values (Fig. 4.16d). As expected, varying β has changed the respective LLT of the EWS. For each sampling proportion, the hesitant-pro-vaccine join count $\langle H, V_s \rangle$ is the best performing join count, matched by $\langle H, N \rangle$ and $\langle N, V_s \rangle$ on the sampled network (Fig. C.3g and C.3h). The two join counts least affected by sampling are $\langle H, V_s \rangle$ and $\langle V_s, V_s \rangle$, both losing less than 20% of their average lead distances when

EWS	β	Penalty (%)			Lead Distance (κ)		
		Minimum	Mean	Maximum	Minimum	Mean	Maximum
$\langle H, H \rangle$	0.4	0	34.38	102.4	0.04	0.59	0.92
	0.6	4.27	38.52	96.42	0.06	0.48	0.86
	0.8	3.7	29.99	81.35	0.14	0.6	0.92
	1	NA	NA	NA	0.51	0.73	0.95
$\langle H, N \rangle$	0.4	35.96	228.31	338.01	0.73	0.86	0.95
	0.6	11.99	207.19	335.61	0.57	0.8	0.95
	0.8	10.39	199.92	313.23	0.61	0.77	0.89
	1	NA	NA	NA	0.04	0.26	0.79
$\langle H, V_s \rangle$	0.4	0	9.01	34.04	0.73	0.86	0.95
	0.6	0	13.33	41.61	0.57	0.8	0.95
	0.8	3.28	10.59	41.1	0.61	0.77	0.89
	1	NA	NA	NA	0.67	0.83	0.98
$\langle N, N \rangle$	0.4	0	40.59	114.35	0	0.04	0.09
	0.6	0	21.29	72.77	0.01	0.06	0.12
	0.8	6.91	63.54	97.04	-0.05	0.02	0.12
	1	NA	NA	NA	0.01	0.06	0.12
$\langle N, V_s \rangle$	0.4	5.24	47.65	147.9	0.73	0.86	0.95
	0.6	0	41.66	142.66	0.57	0.8	0.95
	0.8	4.55	38.39	122.37	0.61	0.77	0.89
	1	NA	NA	NA	0.04	0.6	0.98
$\langle V_s, V_s \rangle$	0.4	0	19.99	82.58	0.06	0.68	0.89
	0.6	0	11.94	50.17	0.42	0.78	0.98
	0.8	3.35	18.88	61.23	0.42	0.68	0.83
	1	NA	NA	NA	0.61	0.81	0.95

Table 4.4: Lead distances and sampling penalties of the join count statistics $\langle *, * \rangle$ for different sampling proportions β , using the SNHT. A green-shaded row indicates that the EWS has mean sampling penalties less than 25% for all sample proportions.

the network was sampled; $\text{Pen} \left[\langle H, V_s \rangle^\beta \right] < 34\%$ and $\text{Pen} \left[\langle V_s, V_s \rangle^\beta \right] < 20\%$ in Tab. 4.4.

4.3.3 Clustering and distance metrics predict the transitions K_*

The final three EWS investigated here are the number of **triads** formed $\langle \Delta_* \rangle$, as well as the **global clustering coefficients** $\langle C_* \rangle$ and the **graph diameters** $\langle \Omega_* \rangle$ of each opinion network (the sub-networks formed by agents holding a specific opinion). The κ -series of each of these EWS are shown in Fig. 4.17b; the panels of the diagram show that some of these EWS are undefined for extreme values of κ . This happens in consensus regimes, where there are simply not enough agents for successful calculation of the statistic. Some of the triad censuses $\langle \Delta_* \rangle$ (Fig. 4.17b) are good indicators of the transitions K_* , with clearly visible pre-transition trends for both values of the social norm $\sigma = 0$ (left) and $\sigma = 0.25$ (right).

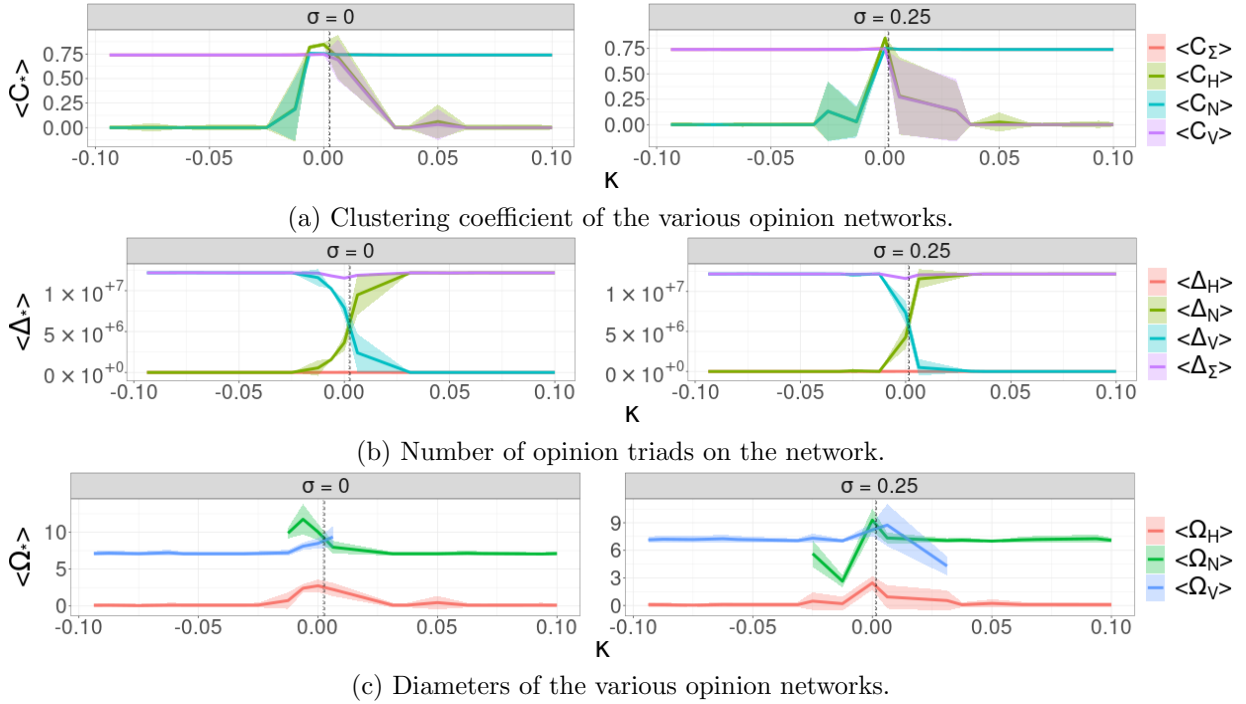


Figure 4.17: Values of the global clustering coefficient $\langle C_* \rangle$, triad census $\langle \Delta_* \rangle$ and opinion network diameter $\langle \Omega_* \rangle$ with respect to the perceived vaccine risk κ . Social norm $\sigma = 0$ for the panels on the left and $\sigma = 0.25$ for the panels on the right.

As was mentioned before, the opinion network diameters are sometimes undefined; for example, the diameter of the N -network $\langle \Omega_N \rangle$ is undefined for smaller extreme perceived vaccine risks of $\kappa \approx -1$ (Fig. 4.17c). This is because $\kappa = -1$ is the pro-vaccine consensus regime, where the small (if nonzero) number of scattered anti-vaccine agents forms a(n) (almost) maximally disconnected network with trivial diameter. The diameter of the H -network is always defined due to the aforementioned equilibrium persistence of hesitant agents for all values of the social norm σ and the grouping also seen in Fig. 4.13c. There is also a corresponding pre-transition rise in the diameter of the H -network as more agents experience hesitance on their way to the vaccine opinion achieving dominance.

Predictably, the increase of the social norm in Fig. 4.17c from $\sigma = 0$ (left) to $\sigma = 0.25$ (right) extends the interval of definition of $\langle \Omega_V \rangle$ and $\langle \Omega_N \rangle$, suggesting that increased social pressure to maintain current opinion prolongs the survival of unpopular opinion (stronger reinforcement from like-minded contacts allows agents to ‘dig in’).

Figure 4.17a shows the trends of the κ -series of the clustering coefficients of the opinion networks.

For both values of the social norm ($\sigma = 0$ for the left panel, $\sigma = 0.25$ for the right panel), the clustering coefficient of the H -network rises and falls symmetrically pre- and post-transition, while the networks supporting the prevalent opinion in consensus regimes predictably have large coefficients. For instance, the clustering coefficient of the V_s -network $\langle C_V \rangle$ is high during the pro-vaccine consensus regime $\kappa \approx -1$ and decreases post-transition. and predictably $\langle C_N \rangle$ shows an opposing pattern.

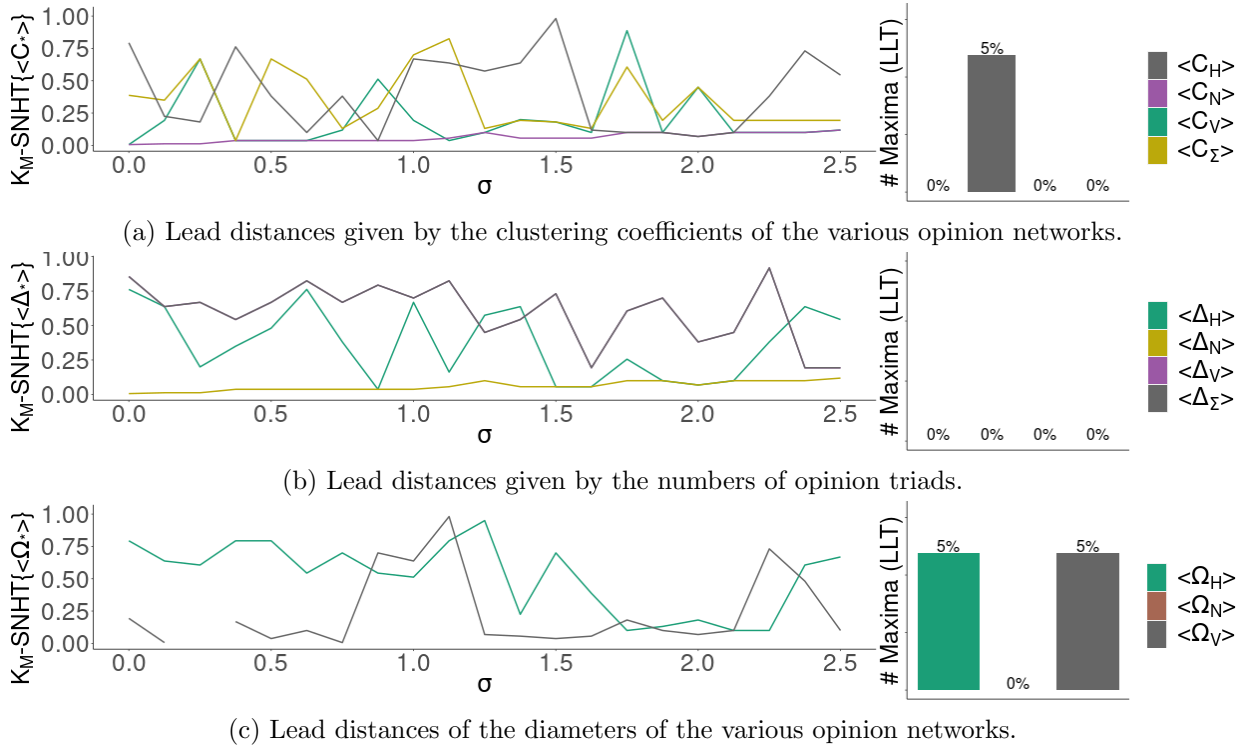


Figure 4.18: Trends in the lead distances given by other connectivity-based EWS investigated using the SNHT.

Surprisingly, these EWS perform relatively badly under the SNHT; an example is the clustering coefficients $\langle C_* \rangle$ (Fig. 4.18). The only EWS that give the largest lead distance for any σ values both measure hesitant groups: $\langle \Omega_H \rangle$ (the diameter of the H network, Fig. 4.18c) and $\langle C_H \rangle$ (the clustering coefficient of the H network, Fig. 4.18a). Lead distances of the number of opinion triads $\text{SNHT}_\sigma\{\langle \Delta_* \rangle\}$ vary wildly (Fig. 4.18b), with the number of anti-vaccine triads $\langle \Delta_N \rangle$ giving consistently low lead distances. The sub-network of hesitant agents gives the highest lead distances $\text{SNHT}_\sigma\{\langle C_H^\beta \rangle\}$ of all clustering coefficients for all sampling proportions β , with LLT 5% with $\beta < 1$ (Figs. C.4a and C.4b). For all proportions β , the clustering coefficient of the anti-vaccine network is the least effective EWS.

Otherwise, the trends in the lead distances of the triad censuses all show high variability with respect to the social norm σ (Figs. C.4c and C.4c). Other than changes to the LLT (resulting from the achievement of a single maximum), changes to the sample proportion β also don't greatly affect the trends of network diameter lead distances $\text{Lead}_\sigma[\langle\Omega\rangle_*^\beta]$ (Figs. C.4e and C.4e). The EWS here least affected by network sampling are the mean number of anti-vaccine triads (closed triangles) on the network ($\langle\Delta_N\rangle$) and the clustering coefficient of the anti-vaccine subnetwork ($\langle C_N\rangle$); both EWS show the same lead distances and sampling penalties in Tab 4.5.

4.3.4 Mutual information, crowned yet again

In sum, we find that mutual information $\langle\mathcal{M}\rangle$ and the number of changes of opinion $\langle\Theta_*\rangle$ (Tab.4.1), the mean sizes of pro-vaccine communities $\langle|Z_V|\rangle$ (Tab. 4.2), the number of communities of hesitant agents $\langle\#Z_H\rangle$ (Tab. 4.3) and the number of interactions among pro-vaccine agents $\langle V_s, V_s\rangle$ and between hesitant and pro-vaccine agents $\langle H, V_s\rangle$ (Tab. 4.4) give similar lead distances despite sampling, unlike other EWS tested. A breakdown of the performance of each EWS with respect to the social norm σ is shown in Fig. 4.19, where the goodness of the lead distance is represented by the colour of the corresponding tile. Green tiles show the location of a *good* warning, which occurs when the EWS gives the largest lead distance of all the EWS; red tiles show when the EWS gave the minimum positive lead distance of all EWS. Black tiles represent failures (negative lead distances), white tiles represent missing or undefined values, yellow tiles indicate social norm σ values for which all EWS gave the same lead distance and grey tiles signify intermediate lead distances (defined values that are neither maximum nor minimum).

Side-by-side bar charts in Fig. 4.20 summarise the grid diagram (Fig. 4.19) by directly showing the ratios of 'good' and 'bad' warnings from each EWS. Specifically, the length of green bars on the right of Fig. 4.20 show the proportion of social norms σ for which the EWS gave *good* warnings (the greatest lead distance of all EWS for that σ value), and so corresponds to the proportion of green tiles in the EWS' row in the grid (Fig. 4.19). Similarly, the length of the red bars on the right of Fig. 4.20 corresponds to the proportion of *bad* warnings, which are minimum positive values (red grid tiles), failures (negative lead distances, black grid tiles) and undefined values (white grid tiles). Figure 4.20 also reveals the mutual information statistic $\langle\mathcal{M}\rangle$ to be the best performer, with no bad warnings and 33% of its warnings being good (i.e., the best of all EWS). The next best EWS were the (minimum, mean and maximum) sizes of pro-vaccine communities $|Z_V|$; however, the number of pro-vaccine communities $\#Z_V$ and others give bad warnings over all σ tested, with negative lead distances at $\sigma = 0.25$ and undefined everywhere else (Fig. 4.19). When sampling 60% of the network ($\beta = 60\%$), the goodness of the mutual information EWS ($\langle\mathcal{M}\rangle$) is matched by the number of communities of hesitant agents Z_H (Fig. C.5b).

Other successful easily-computable EWS are the join counts with pro-vaccine agents: Fig. 4.19 shows that $\langle N, V_s\rangle$ (LLT 14% good), $\langle V_s, V_s\rangle$ (LLT 14%) and $\langle H, V_s\rangle$ (LLT 24%) all perform




Figure 4.19: Grid comparing the performance of each EWS at each value of the social norm σ under the SNHT. Green tiles represent maximum lead distances (of all EWS), grey tiles represent intermediate lead distances, red tiles represent minimum positive lead distances (of all EWS), white tiles represent missing/undefined values, black tiles represent failures (negative lead distances) and yellow tiles indicate that all EWS gave the same lead distance.

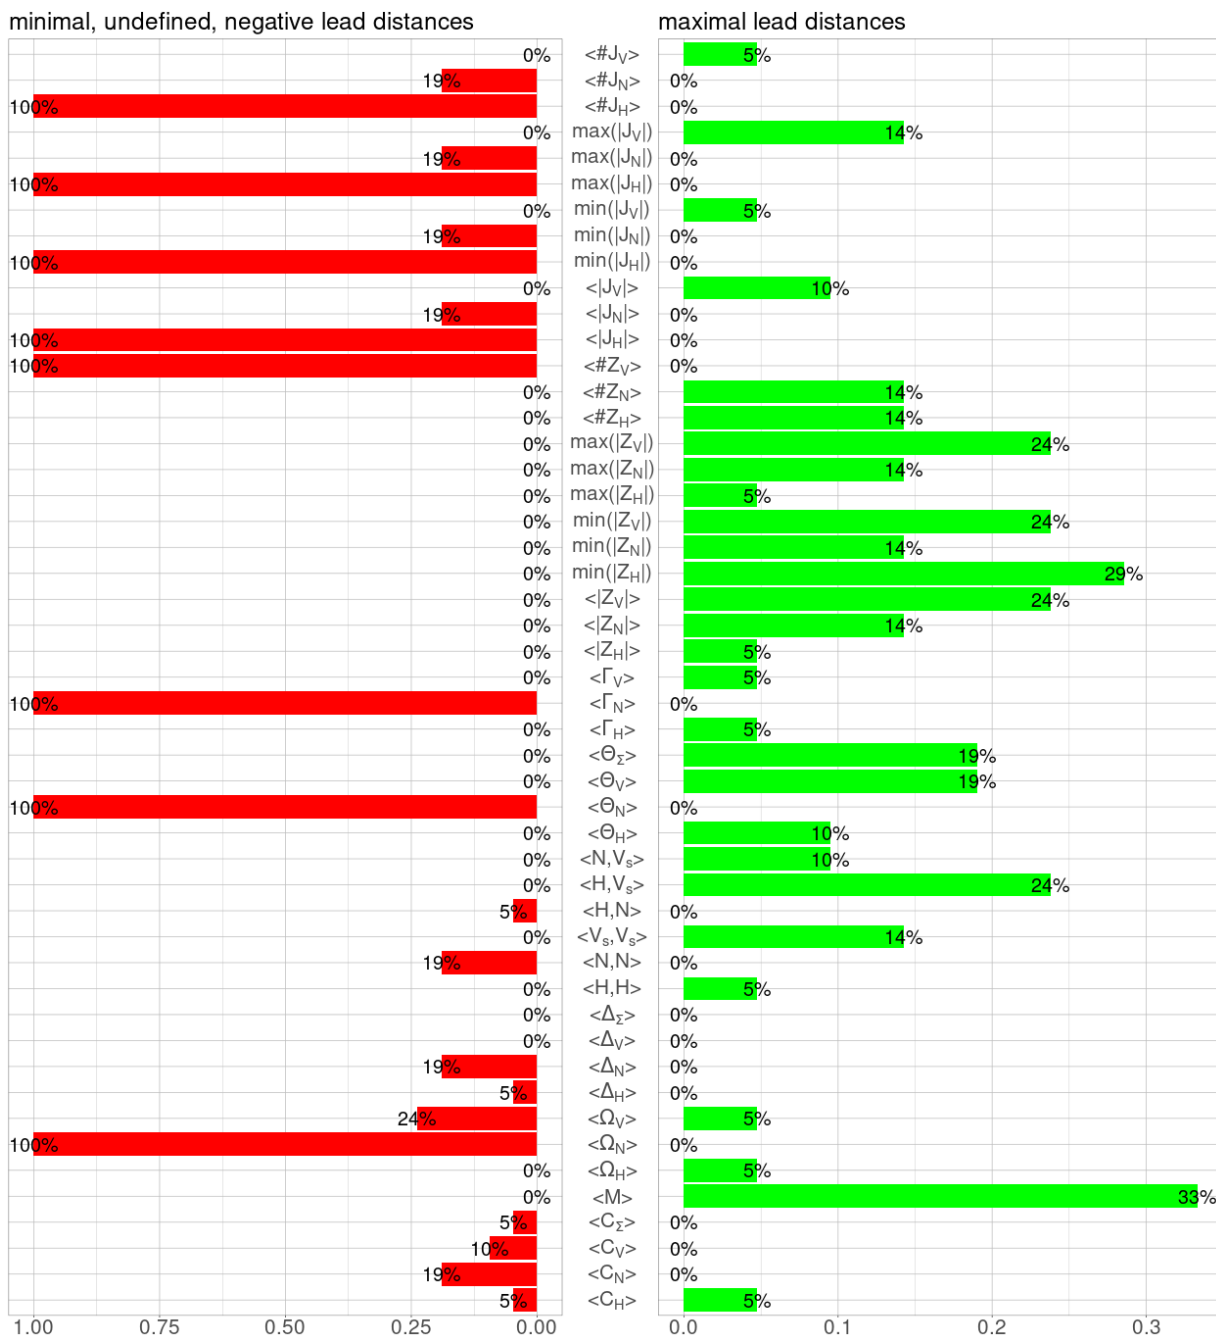


Figure 4.20: Side-by-side bar charts comparing the relative performance of each EWS under the SNHT. Green bars give the proportion of σ values for which the EWS gave the maximum lead distance of all EWS and the red bars give the proportion for which the lead distances were minima, failures or undefined.

well, with no bad warnings given under the SNHT. With a sampling proportion of 80%, the join counts $\langle N, V_s \rangle$, $\langle H, V_s \rangle$ and $\langle H, N \rangle$ take the lead, all with the same LLT of 24%, with $\langle V_s, V_s \rangle$ as fourth-best (with an LLT of 19%). This shows that the reliability of the join counts could be expected. All the EWS tracking echo chambers of hesitant agents J_H yielded all bad warnings for all σ values; this is not surprising given the consistently low number of hesitant agents on the network. Anti-vaccine echo chambers J_N also offered no good warnings under the SNHT; both these results hold for all sample proportions β (Figs. C.5 and C.6). The other three EWS with undefined lead distances for all social norms are all related to the anti-vaccine community; the number of opinion changes $\langle \Theta_N \rangle$, the opinion network diameter $\langle \Omega_N \rangle$ and the probability of having an infected neighbour $\langle \Gamma_N \rangle$ appear as completely white rows in Fig. 4.19. This holds also when the EWS are gathered from random samples of the network, as can be seen in Figs. C.5 and C.6.

4.4 Discussion

In this study, we took early warning signals (EWS) shown to be effective in previous studies of coupled behaviour-infection models and tested their robustness in three ways. By changing the structure of the interaction network, introducing vaccine hesitance to the social dynamics and retrieving the EWS from incomplete samples of the network, we found that most of the EWS used retained their efficacy even in the face of these three challenges. Specifically, the mutual information and join count statistics were overall the best performing EWS, as tested using the standard normal homogeneity test (SNHT).

The introduction of hesitance to the model first required a slightly different method of estimating the location of the social transition. With this, and the change to a small world network, we notice that the two transitions no longer occur in predictable order; infection outbreak actually precedes vaccination crisis on a small area of the parameter space. Also, the distance between the two transitions does not decrease uniformly as the social norm strengthens, as with previous studies. Instead, that gap varies over the parameter space. In Chs. 2 and 3, we interpreted this intertransition distance as allowable reaction time; it seems now that this metaphor is either insufficient, or perhaps that the behaviour of this new model is such that any warning should result in immediate intervention. We realise that such a bleak interpretation as the latter may inspire doubts of the veracity of the model, but the preliminary results shown during parameterisation confirm the correctness of the dynamics and matched other models studied. This also resulted in the re-specification of the lead distance as the distance between the warning given by the EWS and the first transition seen on the network, since we think the EWS are best tested against the worse-case scenario.

This study looked at the re-use of mutual information, join counts, communities, probability of illness, and connectivity metrics as EWS; other tools contraindicated, found ineffective or

redundant in previous studies were excluded to reduce the time taken for simulation and analysis. Parameters chosen during model calibration provided initial run results similar to previous models used, which was considered favourable to model comparison.

Non-connectivity-based EWS such as the mutual information statistic remained effective, showing significant change in trend before the transition and high lead distances during change point testing. There are also many more opinion changes occurring around the transitions in the model dynamics (intermediate values of perceived vaccine risk), which is testament to an active social norm (despite its strength) relative to the perceived vaccine risk. In the anti-vaccine consensus regime, the low number of hesitant individuals were not very likely to have had an infected neighbour; this is intuitive, since the precariousness of the hesitant opinion in this model suggests that an infected neighbour would weigh heavily on the decision to adopt pro-vaccine stance, though this measure ultimately proved not as effective an EWS as the others.

The propagation of opinion communities and echo chambers was also shown to be a strong EWS, in some cases giving high lead distances. As was expected, the shift from pro-vaccine to anti-vaccine consensus saw the breakup and disappearance of pro-vaccine communities and echo chambers, and the establishment and expansion of anti-vaccine groups. Pockets of hesitance appeared around transition, to be quickly eroded in either consensus regime due to the social norm to change opinion. These pockets however were not of sufficient size to feature a significant number of echo chambers of hesitant agents, and indeed the model dynamics discourage this. Join count statistics again showed themselves in a positive light, with clear changes in trend around transitions in the dynamics. The number of joins involving hesitant agents spike around the transitions; due to the static nature of the network, this is possibly simply due to the heightened presence of hesitant agents.

Shown later, the degree of clustering and the triad census present in the corresponding opinion communities follow largely similar patterns of increase and decrease; this is intuitive, since the changes in connectivity quantified by these various metrics are uniform. The one exception is the measurements of the diameters of the pro-vaccine and anti-vaccine groups; the abruptness of the shrinking of these groups appears to depend on the strength of the social norm.

In addition, we established that the EWS are generally resilient to partial information; that is, observation of portions of the network. Lead distances of many of the EWS remain acceptable when even a randomly chosen proportion as small as 60% was gathered, though there are unavoidably performance penalties for omitting almost half the data necessary to establish the initial figures. In sum, we believe that the results of this study both complement and contribute to past and contemporary literature surrounding the behaviour of coupled behaviour-infection models and the use of EWS. Particularly poignant is our establishment of the resilience and reliability of these tools to changes in context and available data, since this instills faith in results and various types of deployment of the tools, as well as potentially allaying any concerns dealing with computational expense or data sparseness, especially concerning a tool as data- and resource-intensive

as mutual information.

We are entirely cognisant of the numerous lingering limitations to our coupled model. These intentional limitations were deemed appropriate to our desire of generating preliminary findings, triaging the pool of EWS candidates, and increasing the complexity of the model only incrementally (a common concern and widely regarded best-practice in our field of computational epidemiology). Some of these limitations are the staticity of the network, the constant rate of neighbour communication and our investigation of the model at equilibrium (rather than focusing on time evolution), all of which are shown to influence model dynamics and the employment of EWS. Despite these, our results constitute a valid contribution to the study of EWS and outbreaks of infectious disease, especially given the recent recurrence of many vaccine-preventable illnesses and the current dire need of coherent frameworks providing warning of vaccine crisis and epidemic events.

EWS	β	Penalty (%)			Lead Distance (κ)		
		Minimum	Mean	Maximum	Minimum	Mean	Maximum
$\langle \Delta_H \rangle$	0.4	3.34	88.85	210.67	0	0.5	0.98
	0.6	0	78.98	188.93	0.04	0.4	0.98
	0.8	2.23	77.55	179.46	0	0.29	0.92
	1	NA	NA	NA	0.04	0.37	0.76
$\langle \Delta_N \rangle$	0.4	0	40.59	114.35	0	0.04	0.09
	0.6	0	21.29	72.77	0.01	0.06	0.12
	0.8	6.91	63.54	97.04	-0.05	0.02	0.12
	1	NA	NA	NA	0.01	0.06	0.12
$\langle \Delta_V \rangle$	0.4	0	35.8	114.15	0.06	0.73	0.98
	0.6	5.14	41.53	98.73	0.06	0.53	0.86
	0.8	0.68	26.8	128.89	0.07	0.48	0.8
	1	NA	NA	NA	0.19	0.61	0.92
$\langle \Delta_\Sigma \rangle$	0.4	0	39.16	143.35	0.06	0.59	0.92
	0.6	5.23	40.91	131.84	0.1	0.5	0.92
	0.8	0.7	29.2	131.14	0.14	0.53	0.8
	1	NA	NA	NA	0.19	0.6	0.92
$\langle \Omega_H \rangle$	0.4	2.29	67.56	168.32	0.01	0.34	0.82
	0.6	17.18	70.01	138.55	0.04	0.38	0.82
	0.8	4.96	54.43	154.96	0.04	0.42	0.92
	1	NA	NA	NA	0.1	0.55	0.95
$\langle \Omega_V \rangle$	0.4	2.73	127.21	421.08	0	0.19	0.95
	0.6	0	96.22	317.18	-0.01	0.2	0.92
	0.8	11.85	132.29	438.41	-0.05	0.16	0.64
	1	NA	NA	NA	-0.01	0.23	0.98
$\langle C_\Sigma \rangle$	0.4	18.17	95.41	256.14	0.04	0.41	0.92
	0.6	0	60.38	174.39	0.04	0.28	0.86
	0.8	10.29	74.13	221.01	0.06	0.3	0.89
	1	NA	NA	NA	0.04	0.34	0.82
$\langle C_H \rangle$	0.4	0	89.56	186.84	0.01	0.49	0.89
	0.6	7.72	79.41	179.12	0.01	0.4	0.95
	0.8	9.78	90.64	232.13	-0.05	0.29	0.92
	1	NA	NA	NA	0.04	0.4	0.98
$\langle C_N \rangle$	0.4	0	40.59	114.35	0	0.04	0.09
	0.6	0	21.29	72.77	0.01	0.06	0.12
	0.8	6.91	63.54	97.04	-0.05	0.02	0.12
	1	NA	NA	NA	0.01	0.06	0.12
$\langle C_V \rangle$	0.4	3.45	137.99	486.99	0	0.19	0.92
	0.6	0	93.42	486.99	-0.01	0.13	0.92
	0.8	9.21	124.89	437.48	-0.05	0.14	0.89
	1	NA	NA	NA	0.01	0.18	0.89

Table 4.5: Lead distances and sampling penalties of the opinion network diameter $\langle \Omega_* \rangle$, triad census $\langle \Delta_* \rangle$ and global clustering coefficient $\langle C_* \rangle$ for different sampling proportions β using the SNHT. A green-shaded row indicates that the EWS has mean sampling penalties less than 25% for all sample proportions.

Chapter 5

Model-based projections for COVID-19 outbreak size and class closures in Ontario

Abstract

The disruption of professional childcare has emerged as a substantial collateral consequence of the public health precautions related to the current COVID-19 pandemic. Increasingly, it is becoming clear that childcare centres must be (at least partially) operational in order to further mitigate the socially debilitating challenges related to pandemic induced closures. However, proposals to safely reopen childcare while limiting COVID-19 outbreaks remain understudied, and there is a pressing need for evidence-based scrutiny of the plans that are being proposed. Thus, in order to support safe childcare reopening procedures, the present study employed an agent-based modelling approach to generate predictions surrounding risk of SARS-CoV-2 transmission and student-days lost within a hypothetical childcare centre. Based on existing proposals for childcare centre and school reopening in Ontario, Canada, six distinct room configurations were evaluated that varied in terms of child-to-educator ratio (15:2, 8:2, 7:3), and family clustering (siblings together vs. random assignment). We also evaluated a primary school setting (30:1, 15:1 and 8:1) including cohorts that alternate weekly. High versus low transmission rates were also contrasted, keeping with the putative benefit of infection control measures within centres, yielding many distinct scenarios. In the childcare scenarios, grouping siblings significantly reduced

Material in this chapter is based on the publication:
Model-based projections for COVID-19 outbreak size and student-days lost to closure in Ontario childcare centres and primary schools. Brendon Phillips, Dillon Browne, Madhur Anand, Chris Bauch. medRxiv 2020.08.07.20170407. <https://doi.org/10.1101/2020.08.07.20170407>.

outbreak size and student-days lost. We identify an intensification cascade specific to classroom outbreaks of respiratory viruses with presymptomatic infection. In both childcare and primary school settings, each doubling of class size from 8 to 15 to 30 more than doubled the outbreak size and student-days lost, by factors of 2-5, respectively 2.5-4.5, depending on the scenario. Across scenarios, having less students per class and grouping siblings together almost always results in significantly lower peaks for number of active infected and infectious cases in the institution. Importantly, the total student-days lost to classroom closure were between 5 and 8 times higher in the 15:2 ratios than for 8:2 or 7:3. These results suggest that current proposals for childcare reopening could be enhanced for safety by considering lower ratios and sibling groupings.

5.1 Introduction

As nations around the world grapple with the psychosocial, civic, and economic ramifications of social distancing guidelines, the critical need for widely-available Early Childhood Education (or colloquially, “childcare”) services have, once again, reached the top of policy agendas [122, 459]. Whether arguments are centred on human capital (i.e., “children benefit from high-quality, licensed educational environments, and have the right to access such care”) or the economy (i.e., “parents need childcare in order to work, and the economy needs workers to thrive”), the conclusion is largely the same: childcare centres are re-opening, at least in some capacity, and this is taking place before a vaccine or herd immunity can mitigate potential spread of SARS-CoV-2 (the virus that causes COVID-19). Outbreaks of COVID-19 in emergency childcare centres and schools have already been observed [515], causing great concern as governments struggle to balance “flattening the curve” and preventing second waves with other pandemic-related sequelae, such as the mental well-being of children and families, access to education and economic disruption.

Governments and childcare providers are tirelessly planning the operations of centres, with great efforts made to follow public health guidelines for reducing SARS-CoV-2 contagion [359]. However, these guidelines, which will result in significantly altered operational configurations of childcare centres and substantial cost increases, have yet to be rigorously examined. Moreover, discussions of childcare are presently eclipsed by general discussion of “school” reopening [114]. That being said, for many parents, the viability of the school day emerges from before- and after-school programming that ensures adequate coverage throughout parents’ work schedules. Yet, reopening plans often fail to mention the critical interplay between school and childcare, even though many childcare centres operate within local schools [498]. Consequently, a model that comprehensively examines the multifaceted considerations surrounding childcare operations may help to inform policy and planning. As such, the purpose of the present investigation is to develop an [agent-based model \(ABM\)](#) that explores and elucidates the multiple interacting factors that could impact potential SARS-CoV-2 spread in school-based childcare centres.

In Ontario, Canada (the authors’ jurisdiction), childcare centres were permitted to reopen on

June 12, 2020, provided centres limit groupings (e.g., classrooms) to a maximum of 10 individuals (educators and children, inclusive) [406]. Additionally, all centres had to come up with a plan for daily screening of incoming persons, thorough cleaning of rooms before and during operations, removal of toys that pose risk of spreading germs, allowing only essential visitors, physical distancing at pick-up and drop-off, and a contingency plan for responding should anyone be exposed to the virus (e.g., closing a classroom or centre for a period of time). Further school-specific recommendations have been recently outlined by The Toronto Hospital for Sick Children [498], which include specific guidelines for screening, hand hygiene, physical distancing, cleaning, ventilation, and masking. While this influential report has become the guiding framework for school reopening in Ontario, there remains no discussion of childcare operations in relation to SARS-CoV-2 spread. Guidelines for primary schools call for either full re-opening, with up to 30 students per classroom attending every day, or with cohorts of 15 students attending in alternate weeks.

Simulation models of infectious disease spread have been widely applied during the COVID-19 pandemic, as in previous pandemics [424, 45]. Modelling is used to determine how quickly the pathogen can spread [428], how easily it may be contained [204], and the relative effectiveness of different containment strategies [327, 190]. Sensitivity analysis is crucial to assess whether model predictions are robust to uncertainties in data [112], which is particularly important during a pandemic caused by a novel emerging pathogen like SARS-CoV-2. ABMs are particularly well-suited to situations where a highly granular description of the population is desirable and where random effects (stochasticity) are important. Such models have been previously applied in both pandemic and non-pandemic situations [306, 299, 555], and such is our choice of modelling methodology in the present work focusing on SARS-CoV-2 transmission in schools and households. Our objective was to use our ABM to project the impact of student-to-educator (or in the case of childcare centres, child-to-educator) ratios and sibling grouping strategies on outbreaks of COVID-19 and student-days lost to classroom closure in a hypothetical childcare centre and primary school.

Below, the modelling approach, results, and interpretation of the present modelling exercise are described. In Sec. 5.2.2, the rationale and parametrisation of the model are specified in detail. In Sec. 5.3, the performance of the model under different assumptions is showcased. We start with analyzing the childcare centre setting and end with the primary school setting. Lastly, the discussion will provide a review and interpretation of this study, including any limitations and future suggestions for research.

5.2 The model

5.2.1 Overview

A detailed description of the model structure, assumptions and Parametrisation appears in Sec. 5.2.2. We developed an ABM of SARS-CoV-2 transmission in a population structured into households and classrooms, as might represent a childcare setting or a small primary school (Fig. 5.1). Individuals were categorised into either child or adult, and contacts between these groups were parameterised based on contact matrices estimated for the Canadian setting.

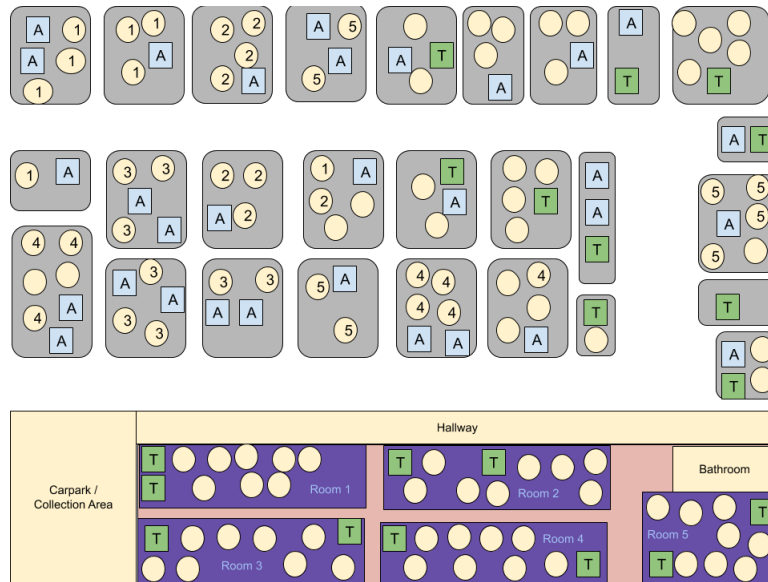


Figure 5.1: Schematic representation of the model population. Blue squares with ‘A’ represent adults, green squares with ‘T’ represent educators, and yellow circles represent children. Grey rectangles represent houses and the school is represented at the bottom of the figure. Numbers exemplify possible assignments of children in households to classrooms.

Household sizes were determined from Canadian demographic data. Classroom sizes and student-educator ratios were determined according to the scenario being studied. For the childcare setting we analysed student-educator ratios of 8:2 and 7:3, giving a maximum class size of 10 representative of the smaller enrolment at schools. We also analysed a student-educator ratio 15:2, giving a total class size of 17. Along with class size, we also consider class composition. Individuals may spread the infection to their household members each day, so effective contacts and interaction in the classroom may result in qualitatively different spreading patterns. As such, children in this model can be assigned to classrooms either randomly (RA) or by grouping siblings (or otherwise

cohabiting students) together (ST) in an attempt to reduce SARS-CoV-2 transmission. For the primary school setting, we considered student-educator ratios of 8:1, 15:1, and 30:1, all with random allocation. For the 8:1 and 15:1 ratios we also considered scenarios where cohorts of 8 or 15 students attending the same classroom but in alternating weeks. These scenarios were labelled 8(A):1 and 15(A):1. In the primary school setting, we considered the higher student-educator ratio 30:1 as an example of larger class size. Some plans considered in reopening Ontario educational institutions divides this larger class size into two alternating cohorts of 15 students each with a single shared educator; we call this scenario 15(A):1. Rotation occurs each week, so that one cohort engages with online material while the other receives face-to-face instruction for 5 days, after which the cohorts exchange roles. The student-educator ratios 8:1 and 8(A):1 were also included for comparison to smaller class sizes. For primary schools we considered only the RA allocation.




Figure 5.2: Diagram showing the SEPAIR infection progression for each agent in the simulation (see Tab. 5.2 for definitions of parameters).

SARS-CoV-2 could be transmitted in households, classrooms or in common areas of the school, all of which were treated as homogeneously mixing on account of evidence for aerosolised routes of transmission [482]. Individuals were also subject to a constant background risk of infection from other sources, such as shopping centres. Figure 5.2 shows the progression of the illness experienced by each individual in the model. In each day, susceptible (S) individuals exposed to the disease via community spread or interaction with infectious individuals (those with disease statuses P , A and I) become exposed (E), while previously exposed agents become presymptomatic (P) with probability δ . Presymptomatic agents develop an infection in each day with probability δ , where they can either become symptomatically infected (I) with probability η or asymptotically infected (A) with probability $1 - \eta$. If a symptomatic individual appears in a classroom, that classroom is closed for 14 days (in the case of alternating cohorts for primary schools, we assumed both cohorts are closed). Other classrooms in the same school may remain open. Asymptomatic students and educators return at the end of this period while symptomatic students and educators

remain at home and symptomatic educators are replaced by substitutes.

Children are less affected by the SARS-CoV-2 virus than adults, and account for a smaller proportion of COVID-19 cases [340]. However, the role of children in SARS-CoV-2 transmission is still debated, and existing epidemiological evidence is limited by lack of empirical studies in school settings, which have been closed for much of 2020. Other studies show that children shed a similar amount of virus to adults [275]. To account for this ambiguity, we used contact matrices drawn from populations under ‘business as usual’ circumstances as a proxy of what contact rates would look like under a full reopening of schools and workplaces [458], but we considered both a high transmission rate scenario and a low transmission rate scenario. The low transmission rate scenario represented either reduced transmission rates in children, and/or highly effective infection control through consistent use of high-effectiveness masks, social distancing, and disinfection protocols (see Sec. 5.2.2 for details). In total the permutations on student-educator ratios, transmission rate assumptions, siblings versus non-sibling groupings, and alternating cohorts yielded 22 scenarios (Table 5.1).

A detailed description of the model structure, assumptions and parametrisation appears in Sec. 5.2.2. We developed an ABM of SARS-CoV-2 transmission in a population structured by households and classrooms, as might represent a childcare setting or a small primary school (Fig. 5.1). Individuals were categorically aged as either child or adult, and contacts between these groups were parametrised based on contact matrices estimated for the Canadian setting. Household sizes were determined from Canadian demographic data. Classroom sizes and student-educator ratios were determined according to the scenario being studied. For the childcare setting we analysed student-educator ratios of 8:2 and 7:3, giving a maximum class size of 10 representative of the smaller enrolment at schools. We also analysed a student-educator ratio 15:2, giving a total class size of 17. Along with class size, we also consider class composition. Individuals may spread the infection to their household members each day, so effective contacts and interaction in the classroom may result in qualitatively different spreading patterns. As such, children in this model can be assigned to classrooms either randomly (RA) or by grouping siblings (or otherwise cohabiting students) together (ST) in an attempt to reduce SARS-CoV-2 transmission. For the primary school setting, we considered student-educator ratios of 8:1, 15:1, and 30:1, all with the random allocation. For the 8:1 and 15:1 ratios we also considered scenarios where cohorts of 8 or 15 students attending the same classroom but in alternating weeks. These scenarios were labelled 8(A):1 and 15(A):1. In the primary school setting, we considered the higher student-educator ratio 30:1 as an example of larger class size. Some plans considered in reopening Ontario educational institutions divides this larger class size into two alternating cohorts of 15 students each with a single shared educator; we call this scenario 15(A):1. Rotation occurs each week, so that one cohort engages with online material while the other receives face-to-face instruction for 5 days, after which the cohorts exchange roles. The student-educator ratios 8:1 and 8(A):1 were also included for comparison to smaller class sizes. For primary schools we considered only the RA allocation.

SARS-CoV-2 can be transmitted in households, classrooms or in common areas of the school, all of which were treated as homogeneously mixing on account of evidence for aerosolised routes of transmission [482]. Individuals were also subject to a constant background risk of infection from other sources, such as shopping centres. Figure 5.2 shows the progression of the illness experienced by each individual in the model. In each day, susceptible (S) individuals exposed to the disease via community spread or interaction with infectious individuals (those with disease statuses P , A and I) become exposed (E), while previously exposed agents become presymptomatic (P) with probability σ . Presymptomatic agents develop an infection each day with probability σ , where they can either become symptomatically infected (I) with probability η or asymptotically infected (A) with probability $1 - \eta$. If a symptomatic case appears in a classroom, that classroom is closed for 14 days (in the case of alternating cohorts for primary schools, we assumed that individuals in both cohorts are kept at home). Other classrooms in the same school remain open. Asymptomatic students and educators return at the end of this period while symptomatic students and educators remain at home and symptomatic educators are replaced by substitutes.

Children are less affected by the SARS-CoV-2 virus than adults, and account for a smaller proportion of COVID-19 cases [340]. However, the role of children in SARS-CoV-2 transmission is still debated, and existing epidemiological evidence is limited by lack of empirical studies in school settings, which have been closed for much of 2020. Other studies show that children shed a similar amount of virus to adults [275]. To account for this ambiguity, we used contact matrices drawn from populations under ‘business as usual’ circumstances as a proxy of what contact rates would look like under a full reopening of schools and workplaces [458], but we considered both a high and a low transmission rate scenario. The low transmission rate scenario represented either reduced transmission rates in children, and/or highly effective infection control through consistent use of high-effectiveness masks, social distancing, and disinfection protocols (see Sec. 5.2.2 for details). In total the permutations on student-educator ratios, transmission rate assumptions, siblings versus non-sibling groupings, and alternating cohorts yielded 22 scenarios (Table 5.1).

5.2.2 Materials and methods

Population structure

There are N households in the population, and a single educational institution (either a school or a school, dependent on scenarios to be introduced later) with M rooms and a maximum capacity dependent on the scenario being tested. Effective contacts between individuals occur within each household, as well as rooms and common areas (entrances, bathrooms, hallways, etc.) of the institution. All groups of individuals (households and rooms) in the model are assumed to be well-mixed.

Each individual (agent) in the model is assigned an age, household, room in the childcare facility and an epidemiological status. Age is categorical, so that every individual is either considered

Childcare centre		
High transmission	15:2 student to educator ratio	siblings together (ST)
		random allocation (RA)
	8:2 student to educator ratio	siblings together
		random allocation
	7:3 student to educator ratio	siblings together
		random allocation
Low transmission	15:2 student to educator ratio	siblings together
		random allocation
	8:2 student to educator ratio	siblings together
		random allocation
	7:3 student to educator ratio	siblings together
		random allocation
Primary school		
High transmission	8:1 student to educator ratio	random allocation
	8:1 student to educator ratio, alternating cohorts	random allocation
	15:1 student to educator ratio	random allocation
	15:1 student to educator ratio, alternating cohorts	random allocation
	30:1 student to educator ratio	random allocation
Low transmission	8:1 student to educator ratio	random allocation
	8:1 student to educator ratio, alternating cohorts	random allocation
	15:1 student to educator ratio	random allocation
	15:1 student to educator ratio, alternating cohorts	random allocation
	30:1 student to educator ratio	random allocation

Table 5.1: Twenty-two scenarios evaluated based on different assumptions about transmission probabilities, educator-student ratios, and student allocation.

a child (C) or an adult (A). Epidemiological status is divided into stages in the progression of the disease; agents can either be susceptible (S), exposed to the disease (E), presymptomatic (an initial asymptomatic infections period P), symptomatically infected (I), asymptotically infected (A) or removed/recovered (R), as shown in Fig. 5.2.

In the model, some children in the population are enrolled as students in the institution and assigned a classroom based on assumed scenarios of classroom occupancy while some adults are assigned educator/caretaker roles in these classroom (again dependent on the occupancy scenario

being tested). Allocations are made such that there is only one educator per household and that children do not attend the same institution as a educator in the household (if there is one), and *vice versa*.

Interaction and disease progression

The basic unit of time of the model is a single day, over which each attendee (of the institution) spends time at both home and at the institution. The first interactions of each day are established within each household, where all members of the household interact with each other. An asymptotically infectious individual of age i will transmit the disease to a susceptible housemate with the age j with probability β_{ij}^H , while symptomatically infectious members will self-isolate (not interact with housemates) for a period of 14 days.

The second set of interpersonal interactions occur within the institution. Individuals (both students and educators) in each room interact with each other, where an infectious individual of age i transmits the disease to some susceptible individual of age j with probability β_{ij}^C . To signify common areas within the building (such as hallways, bathrooms and entrances), each individual will then interact with every other individual in the institution. There, an infectious individual of age j will infect a susceptible individual of age i with probability β_{ij}^O .

To simulate community transmission (for example, public transport, coffee shops and other sources of infection not explicitly modelled here), each susceptible school attendee is infected with probability λ_S . Susceptible individuals not attending the institution in some capacity are infected at rate λ_N , where $\lambda_N > \lambda_S$ to compensate for those consistent effective interactions outside of the institution that are neglected by the model (such as workplace interactions among essential workers and members of the public).

Figure 5.2 shows the progression of the illness experienced by each individual in the model. In each day, susceptible (S) individuals exposed to the disease via community spread or interaction with infectious individuals (those with disease statuses P , A and I) become exposed (E), while previously exposed agents become presymptomatic (P) with probability δ . Presymptomatic agents develop an infection in each day with probability σ , where they can either become symptomatically infected (I) with probability η or asymptotically infected (A) with probability $1 - \eta$.

The capacity of the sole educational institution in the model is divided evenly between 5 rooms, with class size and student-educator ratio governed by one of three basic scenarios: seven students and three educators per room (7:3), eight students and two educators per room (8:2), and fifteen students and two educators per room (15:2). Classroom allocations for children can be either randomised or grouped by household (siblings are put in the same class).

Symptomatically infected agents (I) are removed from the simulation after 1 day (status R) with probability γ_I , upon which they self-isolate for 14 days, and therefore no longer pose a

risk to susceptible individuals. Asymptomatically infected agents (A) remain infectious but are presumed able to maintain regular effective contact with other individuals in the population due to their lack of noticeable symptoms; they recover during this period (status R) with probability γ_A . Disease statuses are updated at the end of each day, after which the cycles of interaction and infection reoccur the next day.

The actions of symptomatic (status I) agents depend on age and role. Individuals that become symptomatic maintain a regular schedule for 1 day following initial infection (including effective interaction within the institution, if attending), after which they serve a mandatory 14-day isolation period at home during which they interaction with no one (including other members of their household). On the second day after the individual's development of symptoms, their infection is considered a disease outbreak centred in their assigned room, triggering the closure of that room for 14 days. All individuals assigned to that room are sent home, where they self-isolate for 14 days due to presumed exposure to the disease. Symptomatically infected children are not replaced, and simply return to their assigned classroom upon recovery. At the time of classroom reopening, any symptomatic educator is replaced by a substitute for the duration of their self-isolation, at the end of which they reprise their previous role in the institution; the selection of a substitute is made under previous constraints on educator selection (one educator per household, with no one chosen from households hosting any children currently enrolled in the institution).

Parametrisation

The parameter values are given in Tab. 5.2. The sizes of households in the simulation were determined from 2016 Statistics Canada census data on the distribution of family sizes [96]. We note that Statistics Canada data only report family sizes of 1, 2 or 3 children: the relative proportions for 3+ children were obtained by assuming that 65% of families of 3+ children had 3 children, 25% had 4 children, 10% had 5 children, and none had more than 5 children. Each educator was assumed to be a member of a household that did not have children attending the school. Again using census data, we assumed that 36% of educators live in homes with no children, where an individual lives alone with probability 0.282, while households hosting 3, 4, 5, 6, and seven adults occur with probability 0.345, 0.152, 0.138, 0.055, 0.021 and 0.009 respectively. Others live with ≥ 1 children in households following the size and composition distribution depending on the number of adults in the household. For single-parent households, a household with a single child occurs with probability 0.169, and households with 2, 3, 4 and 5 children occur with probabilities 0.079, 0.019, 0.007 and 0.003 respectively. With two-parent households, those probabilities become 0.284, 0.307, 0.086, 0.033 and 0.012.

The age-specific transmission rates in households are given by the matrix:

$$\begin{bmatrix} \beta_{1,1}^H & \beta_{1,2}^H \\ \beta_{2,1}^H & \beta_{2,2}^H \end{bmatrix} \equiv \beta^H \begin{bmatrix} c_{1,1}^H & c_{1,2}^H \\ c_{2,1}^H & c_{2,2}^H \end{bmatrix}, \quad (5.1)$$

Parameter	Probability	Baseline Value	Source
η	symptomatic infection	0.6 (adults) 0.4 (children)	TBD TBD
δ	progression, $E \rightarrow P$	0.5/day	[400, 529]
σ	progression, $P \rightarrow I, A$	0.5/day	[400, 529]
γ_I	progression, $I \rightarrow R$	1.0/day	[400, 529]
γ_A	progression, $A \rightarrow R$	0.25/day	[400, 529]
β_{ij}^H	household transmission	0.109	[296], calibrated
β_{ij}^C	classroom transmission	$\beta^C = \alpha_C \beta^H$,	[296], assumption
β_{ij}^O	common area transmission	$\beta^O = \alpha_O \beta^C$,	[296, 458], assumption
ξ	sibling attending same centre	0.8	assumption
Parameter	Meaning	Value	Source
c_{ij}^H	household contact matrix	...	[458]
c_{ij}^C	room contact matrix	...	[458]
		$\alpha_C = 0.75$	
		$\alpha_O = 0.0025$	
λ_i	infection rate due to other sources	1.16×10^{-4} /day	[460], estimated
R_{init}	initial proportion with immunity	0.1	assumption
o	proportion of childless educators	0.36	[96], assumption
	household size distributions		[96]

Table 5.2: Parameter definitions, baseline values and literature sources.

where c_{ij}^H gives the number of contacts per day reported between individuals of ages i and j estimated from data [458] and the baseline transmission rate β^H is calibrated. To estimate c_{ij}^H from the data in Ref. [458], we used the non-physical contacts of age class 0-9 years and 25-44 years of age with themselves and one another in Canadian households. Based on a meta-analysis, the **secondary attack rate** of SARS-CoV-2 appears to be approximately 15% on average in both Asian and Western households [296]. Hence, we calibrated β^H such that a given susceptible person had a 15% chance of being infected by a single infected person in their own household over the duration of their infection averaged across all scenarios tested. As such, age specific transmission is given by the matrix

$$\beta^H \cdot \begin{bmatrix} 0.5378 & 0.3916 \\ 0.3632 & 0.3335 \end{bmatrix}. \quad (5.2)$$

To determine λ_S we used case notification data from Ontario during lockdown, when schools, workplaces, and schools were closed [460]. During this period, Ontario reported approximately 200 cases per day. The Ontario population size is 14.6 million, so this corresponds to a daily infection

probability of 1.37×10^{-5} per person. However, cases are under-ascertained by a significant factor in many countries [310] – we assumed an under-ascertainment factor of 8.45, meaning there are actually 8.45 times more cases than reported in Ontario, giving rise to $\lambda_S = 1.16 \times 10^{-4}$ per day; λ_N was set to $2\lambda_S$.

The age-specific transmission rates in the school rooms are given by the matrix

$$\begin{bmatrix} \beta_{1,1}^C & \beta_{1,2}^C \\ \beta_{2,1}^C & \beta_{2,2}^C \end{bmatrix} \equiv \beta^C \begin{bmatrix} c_{1,1}^C & c_{1,2}^C \\ c_{2,1}^C & c_{2,2}^C \end{bmatrix} \equiv \beta^C \begin{bmatrix} 1.2356 & 0.0588 \\ 0.1176 & 0.0451 \end{bmatrix}, \quad (5.3)$$

where c_{ij}^C is the number of contacts per day reported between age i and j estimated from data [458]. To estimate c_{ij}^C from the data in [458], we used the non-physical contacts of age classes 0–9 years and 20–54 years of age, with themselves and one another, in Canadian schools. Epidemiological data on [secondary attack rates](#) in childcare settings are rare, since schools and schools were closed early in the outbreak in most areas. We note that contacts in families are qualitatively similar in nature and duration to contacts in schools with small group sizes, although contacts are generally more dispersed among the larger groups in rooms than among the smaller groups in households. On the other hand, rooms may represent equally favourable conditions for aerosol transmission, as opposed to close contact. Hence, we assumed that $\beta^C = \alpha_C \beta^H$, with a baseline value of $\alpha_C = 0.75$ based on more dispersed contacts expected in the larger room group, although we varied this assumption in sensitivity analysis.

To determine β^O we assumed that $\beta^O = \alpha_O \beta^C$ where $\alpha_O \ll 1$ to account for the fact that students spend less time in common areas than in their rooms. To estimate α_O , we note that β^O is the probability that a given infected person transmits the infection to a given susceptible person. If students and staff have a probability p per hour of visiting a common area, then their chance of meeting a given other student/staff in the same area in that area is p^2 . We assumed that $p = 0.05$ and thus $\alpha_O = 0.0025$. The age-specific contact matrix for β^O was the same as that used for β^C (5.3).

Model Initialisation

Upon population generation, each agent is initially susceptible (S). Individuals are assigned to households as described previously, and children are assigned to rooms either randomly or by household. We assume that parents in households with more than one child will decide to enrol their children in the same institution for convenience with probability $\xi = 80\%$, so that each additional child in multi-child households will have probability $1 - \xi$ of not being assigned to the institution being modelled.

Households hosting educators are generated separately. As above, we assume that 36% of educators live in adult-only houses, while the other educators live in houses with children, both household sizes following the distributions outlined in Sec. 5.2.2. The number of educator households

is twice that required to fully supply the school due to the replacement process for symptomatic educators outlined in Sec. 5.2.2. Initially, a proportion of all susceptible agents R_{init} is marked as removed/recovered (R) to account for immunity caused by previous infection moving through the population. A single randomly chosen school attendee is chosen as a primary case and is made presymptomatic (P) to introduce a source of infection to the model. All simulations are run until there are no more potentially infectious (E, P, I, A) individuals left in the population and the institution is at full capacity. All results were averaged over 2000 trials.

Estimating β^H

Agents in the simulation were divided into two classes: “children” (ages 0 – 9) and “adults” (ages 25 – 44). Available data on contact rates[458] was stratified into age categories of width 5 years starting at age 0 (0 – 5, 5 – 9, 10 – 14, etc.). The mean number of contacts per day c_{ij}^H for each class we considered (shown in (5.2)) was estimated by taking the mean of the contact rates of all age classes fitting within our presumed age ranges for children and adults.




Figure 5.3: Plot showing the probability of infection stemming from single infection in the household with respect to the value of the contact rate coefficient β^H . The shaded region represents one standard deviation of ensemble values obtained for each value of β^H .

For β^H calibration, we created populations by generating a sufficient number of households to fill the institution in each of the three tested scenarios; 15:2, 8:2 and 7:3. In each household, a single randomly chosen individual was infected (each member with equal probability) by assigning them a presymptomatic disease status P ; all other members were marked as susceptible (disease status S). In each day of the simulation, each member of each household was allowed to interact with the infected member, becoming exposed to the disease with probability given in (5.2). Upon exposure, they were assigned disease status E . At the beginning of each subsequent day, presymptomatic

individuals proceeded to infected statuses I and A , and infected agents were allowed to recover as dictated by Fig. 5.2 and Tab. 5.2. This cycle of interaction and recovery within each household was allowed to continue until all infected individuals were recovered from illness.

We did not allow exposed agents (status E) to progress to an infectious stage (I or A) since we were interested in finding out how many infections within the household would result *from a single infected household member*, as opposed to added secondary infections in later days. At the end of each trial, the specific probability of infection (π_n) in each household H_n was calculated by dividing the number of exposed agents in the household (E_n) by the size of the household $|H_n|$ less 1 (accounting for the member initially infected). Single occupant households ($|H_n| = 1$) were excluded from the calculation. The total probability of infection π was then taken as the mean of all π_n , so that

$$\pi = \frac{1}{D} \sum_n \pi_n = \frac{1}{D} \sum_{|H_n| \geq 2} \frac{E_n}{|H_n| - 1}, \quad (5.4)$$

where D represents the total number of multiple occupancy households in the simulation. This modified disease simulation was run for 2000 trials each of different prospective values of β^H ranging from 0 to 0.21. The means of all corresponding final estimates of the infection rate were taken per value of β^H , and the value corresponding to a infection rate of 15% was interpolated as shown in Fig. 5.3.

5.3 Results

5.3.1 Initial stages of the outbreak

The time evolution of the outbreaks are illustrated in Fig. 5.4, which shows the proportion of actively infected school attendees (both children and educators) per day in twelve childcare centres scenarios. Many of the scenarios tend to produce a well-defined outbreak curve close to the start of the simulation, even with classroom closure protocols in place. However, the outbreaks are more strongly household-driven for the 7:3 and 8:2 ratios than the 15:2 ratio; this is apparent in the weekly waves superimposed on the overall epidemic curve more strongly in the 15:2 scenarios, on account of the impact of weekends. The 15:2 ratio also tends to generate earlier, more intense outbreaks, while 7:3 and 8:2 scenarios produce fewer infections that are more sporadically distributed throughout the simulated time horizon. In the case of high transmission, the maximum mean level of exposure (E) is 4.97% in the 15:2 RA configuration 18 days into the the simulation, on average, with peak 3.03% presymptomatic (P) and 1.64% asymptomatic (A) attendees at days 12 and 19 respectively. Meanwhile, peak mean exposure in scenario 7:3 ST occurs on day 2, with 1.9% attendees exposed to the disease, with presymptomatic cases never exceeding that of the start of any simulation.

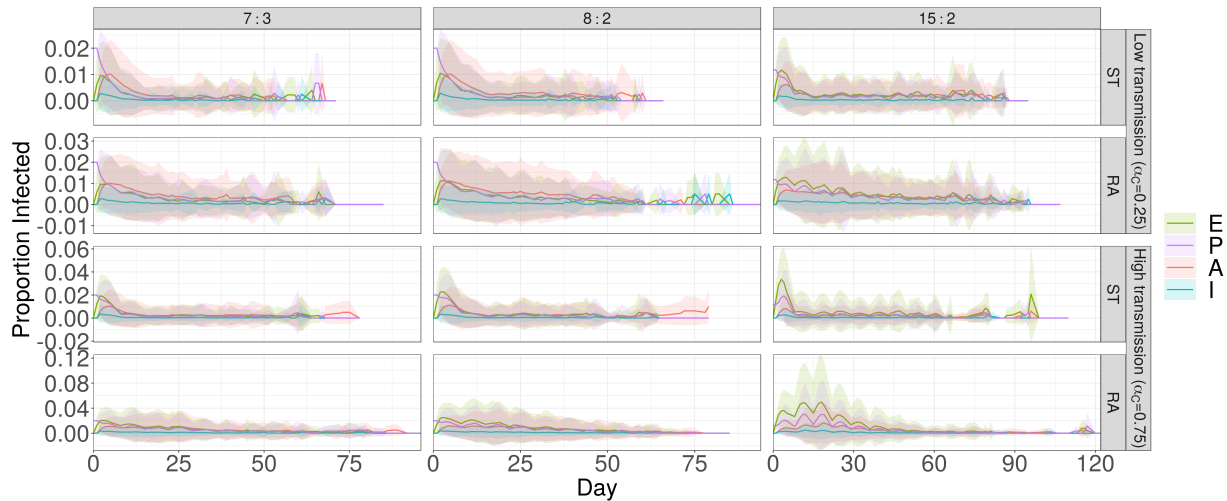


Figure 5.4: Time series of the proportions of exposed (E), presymptomatic (P), asymptomatic (A) and infected (I) individuals in the simulation for each reopening scenario. The ensemble means are represented by solid lines, while the respected shaded ribbons show one standard deviation of the results.

Table 5.3 summarises the information from the figures, showing the days until the 30-day peak of each proportion of active infections in the centres. Here we can see that active infections peak far earlier with the ST allocation than with the RA allocation for both high ($\alpha = 0.75$) and low ($\alpha = 0.25$) transmission rates in most cases, and have either equal or smaller peaks for most maximum proportions corresponding to the RA allocation independent of student-educator ratio. In the case of high transmission, peak proportions decrease with the number of students per class in half of the tested scenarios (statuses P and I with RA allocation, and status E). In the low transmission case, there is a reversal in trend, with peak proportions increasing with decreasing number of exposed (E) and presymptomatic (P) students per class. There is no obvious relationship between peak days for infected (I) and asymptomatic (A) individuals in the high transmission case, nor for asymptomatic (A) individuals in the low transmission case. In all cases (save status A in the low transmission scenario and statuses P , E and I in the high transmission scenario, all with ST allocation), peak proportions decreased consistently with the number of students per classroom. In sum, having fewer students per classroom and grouping siblings together almost always results significantly lower peaks number of active infected and infectious cases in the school. Peaks may also occur sooner in the ST allocation. This may reflect household members spending more time together than under the RA allocation, resulting in a more rapid start to the outbreak even if the number of peak cases is more restricted under the ST allocation.

α_C	Status	Allocation	Peak Time			Maximum ($\times 10^{-4}$)		
			15:2	8:2	7:3	15:2	8:2	7:3
0.75	P	RA	12	0	0	304	200	200
		ST	4	0	0	193	200	199
	E	RA	18	3	3	497	252	204
		ST	3	3	2	336	227	195
	I	RA	12	2	2	49	37	35
		ST	4	2	2	30	34	37
	A	RA	19	5	5	165	198	111
		ST	5	5	4	82	113	103
α_C	Status	Allocation	15:2	8:2	7:3	15:2	8:2	7:3
0.25	P	RA	0	0	0	118	200	200
		ST	0	0	0	118	200	201
	E	RA	4	3	5	96	113	128
		ST	2	2	3	96	105	117
	I	RA	2	2	2	19	27	21
		ST	2	2	2	19	30	21
	A	RA	5	4	5	69	111	100
		ST	5	5	5	62	102	102

Table 5.3: Times at which the mean proportions of presymptomatic (P), exposed (E), symptomatically infected (I) and asymptotically infected (A) school attendees peak during the first 30 days of simulation with secondary spread with respect to each of the scenarios tested, and the corresponding peak number of cases.

The [basic reproductive ratio](#) R_0 is the average number of secondary infections produced by a single infected person in an otherwise susceptible population [20]. When there is pre-existing immunity, as we suppose here, we study the [effective reproduction number](#) \mathcal{R}_e - the average number of secondary infections produced by a single infected person in a population with some pre-existing immunity. Figure 5.5 shows the estimated \mathcal{R}_e and mean population size (school plus all associated households) over the course of each simulation, computed by tracking the number of secondary infections produced by a single primary case. The \mathcal{R}_e values measured from the simulation range from 1.5 to 3 on average, depending on the scenario. These \mathcal{R}_e values are generally lower than the typical range of R_0 values between 2 to 3 reported in the literature [335]. This is the expected relationship, not only because of pre-existing immunity, but also because the \mathcal{R}_e values in our simulation capture transmission only in schools and workplaces, while the R_0 values in the literature are measured for SARS-CoV-2 transmission in all settings, including workplaces and other sources of community spread.

There is little correlation between mean population size (Fig. 5.5, line), number of households




Figure 5.5: Bar chart showing the effective reproduction number \mathcal{R}_e in the entire population (with error bars denoting one standard deviation) and a line plot showing the mean population size. Both low and high transmission scenarios are shown.

(not shown) and the corresponding \mathcal{R}_e estimate (Fig. 5.5, bars), leaving only the number of children per classroom responsible for the gross increasing trend in \mathcal{R}_e in both high ($\alpha = 0.75$) and low ($\alpha = 0.25$) transmission scenarios. Equation 5.3 shows that child-child contact within the classroom occurs at least 2 times more often than any other type of contact; given that the majority of the attendees of the school are children, we can expect \mathcal{R}_e to depend on the number of children enrolled in the school.




Figure 5.6: Diagram showing the proportion of trials without secondary spread (curve), and the time taken to produce the first secondary infection (bar chart), both sorted by scenario.

This is further demonstrated by the bar charts of Fig. 5.6, which show the distribution of times between the primary infection case and the first secondary infection. The scenarios with the highest ratio of children to educators (15:2) show the quickest start of the outbreak in both

high and low transmission cases, with RA having the highest proportion of trials where the first secondary infection occurred within a single day in the high transmission case. In comparison, scenario 7:3 RA showed the slowest average initial spread in the high transmission case, while the low transmission case sees low rates for both 8:2 and 7:3. Configuration ST (except for ratio 7:3) frequently results in faster secondary spread over the first two days (even in the first 2 weeks).

5.3.2 Outbreak duration

Each individual simulation ends when all classes are at full capacity and there are no active infections in the population; aside from community infection, this marks the momentary halt of SARS-CoV-2 spread. From this, we get a description of the duration of the first outbreak. (There could well be a second outbreak sparked by some community infection among individuals who remain susceptible at the end of the first outbreak). Box plots in Fig. 5.7 show that the 15:2 ratio in both RA and ST allocations gives a median outbreak duration at least as large as all other scenarios (for both low and high transmission cases). Another general observation is that classroom allocation (RA vs. ST) doesn't change the distribution of outbreak duration for student-educator ratios 8:2 and 7:3 as drastically as it does for 15:2, whereas ST allocation results in lower median duration (24 vs 43 for RA allocation) and significantly lower maximum duration for the 15:2 ratio (61 vs. 88 for RA allocation without outliers) in the high transmission case.




Figure 5.7: Box plots depicting the distribution of simulation durations for each scenario. Taken together with the stopping criteria of the simulations, these describe the duration of the outbreak. Red dots represent the arithmetic mean of the data.

This is mirrored in the low transmission case as well. A possible explanation lies in the number of students per classroom. The child-child contact rate (5.3) is far higher than any other contact rate, implying that the classroom is the site of greatest infection spread (demonstrated in Fig. 5.9). ST allocation differs from RA allocation in its containment of disease transfer from the classroom

to a comparatively limited number of households. This effect (the difference between ST and RA) is amplified with the addition of each new student to the classroom, so that while the difference between 7:3 and 8:2 may be small (only 1 student added), the effect becomes far exaggerated when the student number is effectively doubled (15 students vs. 7 or 8).




Figure 5.8: Time series detailing the trends in the mean proportions of current school attendees in each stage of disease progression. Shaded ribbons around each curve show one standard deviation of the averaged time series. Only trials showing secondary spread were included in the ensemble means shown.

The evolution of the numbers of susceptible (S) and recovered/removed (R) school attendees provides additional information on the course of the outbreak, since they represent the terminal states of the infection process in each individual by the end of the outbreak. Figure 5.8 shows the proportion of susceptible and recovered current school attendees (who have not been sent home due to classroom outbreaks). As with all results so far, the 15:2 RA scenario most efficiently facilitates disease spread through the school in both high and low transmission cases, with the proportion of recovered attendees (R) overtaking the number of never-infected attendees (status S) on day 34 in the case of high transmission ($\alpha = 0.75$). Performance between 8:2 and 7:3 with ST allocation is similar for both transmission rates, though all scenarios show smaller variation over trials featuring lower infection transmission. As shown in Fig. 5.7, scenario 15:2 RA gave the longest average simulation time in the high transmission scenario; this is also reflected in Fig. 5.4, where the longest outbreak lasted 134 days.

5.3.3 Outbreak size and classroom closure

Figure 5.9 shows the mean number of infections in each location in all scenarios, as well as the total number of infections in each scenario (the ‘outbreak size’). As expected, many more infections occur in the high transmission scenario ($\alpha = 0.75$), and the error bars of the plot show greater standard deviation of the results than in the low transmission ($\alpha = 0.25$) scenario. But for each location and regardless of the transmission rate scenario, the number of infections increases




Figure 5.9: The mean number of infections occurring among all school attendees in each location over time for each scenario. The height of each bar gives the ensemble mean and its standard deviation is represented by error bars.

rapidly with the number of children in the classroom in each room allocation. The 15:2 ratio is universally the worst allocation across all possible scenarios. However, the difference between the outbreak size in different scenarios decreases as the transmissibility of the virus drops (so to speak, the gap between the 15:2 RA and 15:2 ST scenarios decreases as α decreases, and so with other student-educator ratios). When the transmission rate is high, the relatively larger variety (by household) and prevalence of child-child interactions has a multiplicative effect on the number of effective transmissions in the classroom. Lower transmissibility thereby decreases the classroom infection rates relative to the household transmission rates.




Figure 5.10: Box plots showing the number of student days lost over the course of the simulation due to class closure upon the detection of an outbreak. Red text boxes show the mean and standard deviation of total closure durations.

The numbers of student-days lost due to classroom closure are given in Fig. 5.10, according to

scenario. (The number of student-days lost is the number of days of closure times the number of students who would otherwise have been able to continue attending.) In all scenarios, the 15:2 student-educator ratio is quantitatively the worst strategy examined by almost an order of magnitude, resulting in the highest possible number of student-days lost. RA allocation shows worse performance than ST in all scenarios. Both the low ($\alpha = 0.25$) and high ($\alpha = 0.75$) transmissibility scenarios favour the 7:3 student-educator ratio and ST allocation, with a lower number of student-days lost. The poor performance of 15:2 ratio occurs because it suffers from a multiplicative effect: larger class sizes are more likely to be the origin of outbreak, and when the outbreak starts, more children are affected when the classroom is shut down. Moreover, since it's possible for a student or educator to be infected during a 14-day closure, not all attendees necessarily return to class upon reopening; sick educators are replaced with substitutes. As such, these class closures results in otherwise healthy students missing potentially additional school days beyond the 14-day closure period. The 15:2 strategy suffers particularly from this effect, since transmission is facilitated when more students are in a classroom.




Figure 5.11: Bar chart showing the number of days for which some number of rooms in the school were closed due to disease outbreak. Scenarios are represented by different colours; the height of each bar gives the relevant ensemble mean with its standard deviation represented by error bars.

Naturally, a high incidence of COVID-19 cases will result in multiple room closures; one way to see this is to look at the number and duration of room closures, both shown in Fig. 5.11. In all scenarios, schools spent (on average) more days with one closed classroom than any other number. We can also observe a difference in RA and ST allocations for the 7:3 ratio: with both high and low transmission rate ($\alpha = 0.25$ and $\alpha = 0.75$ respectively), RA allocation results in a higher number of class closures.

5.3.4 Primary school settings

The primary school setting shows the same cascade of intensifying outbreaks and rapidly mounting student-days of closure as class sizes increase (Fig. 5.12). This effect occurs in both childcare centres and primary schools because firstly, in a larger classroom it is more likely that a student tests positive for COVID-19. Secondly, when the classroom closes as a result, more students are affected by the closure. Thirdly, because COVID-19 is characterised by presymptomatic infection and aerosol dispersal, there is more infection in larger classrooms before the closure is enacted. Introducing more children into the classroom increases the effective reproductive ratio (\mathcal{R}_e) for both low and high rates of transmission, while cohorting/alternation has little effect (Fig. 5.12C), and similar strategies (that is, differing by only 1 student or educator per class, or by alternation) give similar reproductive ratios \mathcal{R}_e (compare to Fig. 5.5).

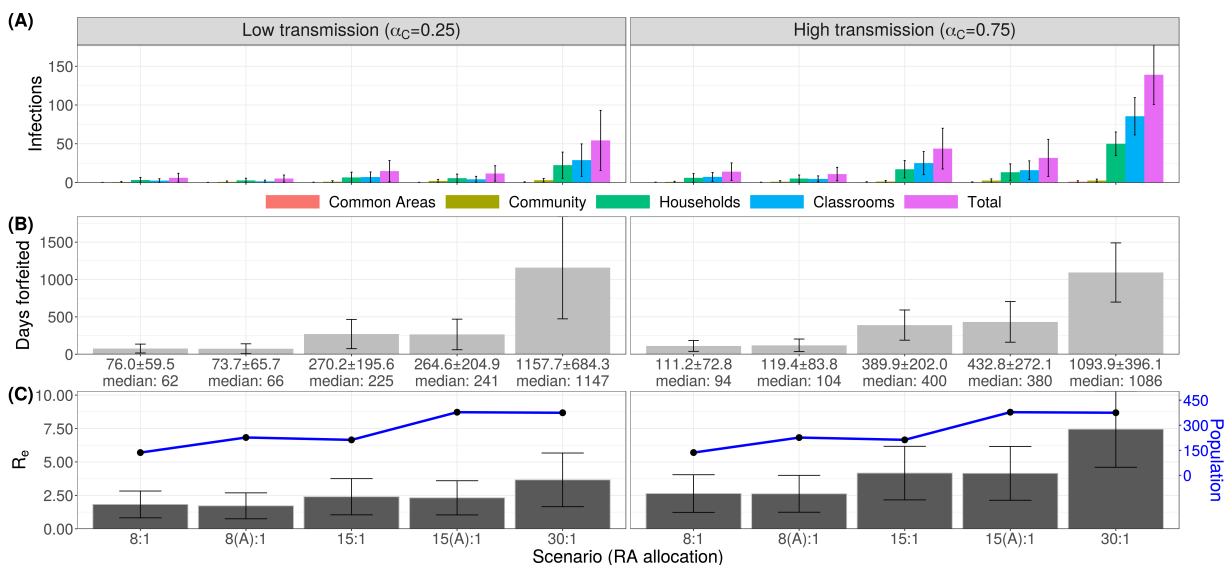


Figure 5.12: Further effects of varying class size when cohorting is introduced. (A) Bar chart showing the total number of infections produced in each location in the model over the time of the simulation. (B) Bar chart showing the number of student days lost due to outbreak-induced classroom shutdowns. (C) Bar chart showing the effective reproduction number \mathcal{R}_e in the entire population, with a line plot showing the mean population size. Both low and high transmission scenarios are shown. In each panel, error bars denote a single standard deviation of the data plotted.

There is little difference between numbers of lost student days between the similar scenarios 8:1 and 8(A):1, as well as as 15:1 and 15(A):1 (Fig. 5.12B). Since the shutdown of a classroom affects both cohorts, there will be very little difference in virus spread between scenarios allotting

the same number of students per classroom. This effect is also seen in Fig. 5.12A. Comparison of Fig. 5.9 and Fig. 5.12A show similar distributions of outbreak size for all student-teacher ratios, signifying that cohorting does not significantly change the results of structured interactions featured in the model. The true benefit of cohorting arises in the consideration of class sizes, given the desire for contact time with all enrolled students. Comparison of Fig. 5.10 and Fig. 5.12B shows that the similar scenarios 15:2 RA, 15:1 RA and 15(A):1 RA all result in a comparable number of lost student-days in both low and high transmission scenarios, as do the scenarios 8:2 RA, 8:1 RA and 8(A):1 RA.

Higher student-educator ratios facilitate faster disease spread through the school than smaller ones (Fig. 5.13). One major difference is the weekly fluctuation of the infection status curves visible in the cohorted scenarios 8(A):1 and 15(A):1. These fluctuations correspond to the rotation of the student cohorts through the school term. Transitions between majority susceptible and recovered regimes is delayed (high transmission) or prevented (low transmission) by cohorting; we see that alternating strategies result in better aggregate infection outcomes, even when classroom capacity is held constant. Scenario 15(A):1 also results in shorter mean and median outbreak lengths in the entire population in both low and high transmission cases (Fig. 5.14).




Figure 5.13: Time series showing the trends in the mean proportions of current school attendees in each stage of disease progression. Shaded ribbons around each curve show one standard deviation of the averaged time series. Only trials showing secondary disease spread were included in the ensemble means shown.

5.3.5 Sensitivity Analysis

We conducted a sensitivity analysis on β^H , β^C , λ and R_{init} (see D). We found that variation in rates of household and classroom interaction and infection (β^H and β^C) and the number of individuals initially recovered (R_{init}) greatly impact SARS-CoV-2 transmission, but did not




Figure 5.14: Box plots depicting the distribution of simulation durations for each scenario, describing the length of the outbreak. Red dots represent the arithmetic mean of the data.

change the relative performances of the 22 scenarios. The greatest influence on outcomes remain the scheme of allocation of students to classrooms (RA or ST), the number of students per class (15, 8 or 7), and whether the transmission rate in the classrooms is low or high (α_C). Other important factors include classroom closure upon identification of a symptomatic case and the interaction patterns of asymptomatic infected individuals in the household upon classroom closure (i.e. whether they continue to interact in close contact, as would be necessary for younger children, or whether children are old enough to effectively self-isolate). Our baseline assumption was to assume asymptomatic infected individuals who are sent home due to closure of a classroom are able to self-isolate. This assumption is conservative, since inability to self-isolate under these circumstances would result in higher projected outbreak sizes.

5.4 Discussion

We developed and simulated an ABM of SARS-CoV-2 transmission in childcare centre and primary school settings for the purposes of informing reopening policies. The model was configured to capture SARS-CoV-2 transmission in a local school building, since many childcare centres operate across several classrooms within schools. These services are an essential bridge for many parents who are unable to drop off or pick up children around school hours due to work. Our findings suggest that variability in class size (i.e., number of children in a class) and class composition (i.e., sibling groupings versus random assignment) influence the nature of SARS-CoV-2 transmission within the childcare context. Specifically, a 7:3 student-to-educator ratio that used sibling groupings yielded the lowest rates of transmission, while a 15:2 ratio consistently performed far worse. Findings for the primary school ratios show a similar acceleration of negative impacts with increasing class size. Findings from our simulations are sobering, as educators in the province lobbied for a 15 student cap on classrooms in Summer 2020. Our study suggests

that classes of this size pose a tangible risk for COVID-19 outbreaks, and that lower ratios would better offset infection and school closures. While school reopening guidelines [498], public health agencies [396], and public petitions [514] have called for smaller class sizes, governments appear to be following some recommendations in reopening plans while ignoring others.

This accelerating effect of increasing classroom sizes occurs because of three factors working in concert. Firstly, a larger class means that a student is more likely to test positive for COVID-19 at some point. Secondly, when a larger class is closed as a result, it affects more students. Third, presymptomatic transmission and higher densities of students ensure that more children become infected before classroom closure is enacted, resulting in larger outbreak sizes due to more cases both before the closure, and after the closure as the infection continues to spread in households. This particular mechanism is specific to institutional outbreaks for infectious diseases with presymptomatic transmission worsened by aerosol transmission routes [482].

Policies related to childcare and traditional school reopening have not been well integrated [160]. In Ontario, childcare classrooms were capped at a maximum of 10 occupants, overall (hence the 8:2 and 7:3 ratios in the present study) [406]. Conversely, procedures for traditional “school” classrooms have been given the go-ahead for 15 children (hence the 15:2 ratio). While allowable class sizes will differ somewhat as a function of child age and jurisdiction, it seems likely that early childhood and elementary school classes may actually surpass these numbers in Ontario. Our findings demonstrate that the 15:2 ratio represents a significantly higher risk, not only for SARS-CoV-2 spread, but for school closures. In one scenario (15:2 random assignment), the modelled outbreak lasted for 105 days. Given that childcare and schools are often operating within the same physical location, this policy discrepancy is questionable. Based on our simulations, a lower ratio (7:3) is indicated. Moreover, it appears that this configuration could be enhanced through the utilization of sibling groupings.

Our modelling approach was informative in terms of identifying the location of SARS-CoV-2 transmission. There has been conflicting evidence on classroom based transmission of SARS-CoV-2 [250, 495]. The present study suggests that classrooms and households yield much higher rates of infection than common areas. Thus, initiatives to reduce inter-classroom contact in common areas (such as staggering start times, utilizing multiple entrances, and sanitizing surfaces in building foyers) may only produce a modest benefit for reducing spread. Conversely, our simulations demonstrated a marked benefit associated with lower child-to-educator ratios in classrooms. Notably, these benefits were observed in both high transmission settings (e.g., at the start of the pandemic, before social distancing), and low transmission settings (e.g., where masking, hygiene, and social distancing had been put in place, as will be the case in reopened childcare and school). Other investigators have proposed intermittent occupancy and enhanced ventilation as potential measures for reducing classroom (indoor) transmission amongst children [358].

An examination of student days missed due to classroom closure further elucidates the favourability of smaller class size and sibling grouping as a preventative measure. In this analysis, the

worst configuration was the 15:2 random assignment ratio. Again, this was observed in both high transmission and low transmission environments. In the most unfavourable scenario (15:2 RA), there were cumulatively 387 and 267 student days lost in high versus low transmission settings, respectively. Conversely, in the best scenario (7:3, siblings together), there were only 47 and 40 student days lost. Thus, our simulations suggest that the lower ratios and sibling groupings offer a safeguard against highly disruptive classroom closures [50, 546]. Given this, a proactive and preventative approach incorporating realistic levels of reduced class time would be better than a reactive strategy that yields unpredictable closure events due to outbreaks.

Several policy and procedural recommendations have emerged from this modelling exercise. First, it is recommended that childcare and school settings, alike, consider lowering student-to-educator ratios. Commensurate with the present findings, a 7:3 ratio (10 individuals per class including both children and adults) outperforms a 15:2 ratio on key metrics. Second, there also appears to be benefit associated with sibling groupings. Thus, a siblings-together configuration should be considered. Third, the majority of transmission occurred in the classroom. As such, it is important for reopening plans to consider social distancing and hygiene procedures within classrooms - a recommendation that may only be feasible with fewer children in the classroom. It is unlikely that classrooms with 15 or more children will afford children the necessary space to socially distance. Finally, in the primary school setting, significant benefits accrue for 15(A):1 relative to the 30:1 student-educator ratio, and thus decision-makers should reconsider the conventional model of putting 30 students in classrooms every day in favour of cohorts of 15 students alternating weekly.

Finally, the present study has a number of limitations that should be considered. While it is becoming increasingly clear that COVID-19 risk varies as a function of social determinants of health (e.g., socioeconomic status, race, ethnicity, immigration status, neighbourhood risk), along with opportunities for social distancing [503], the present study did not take these considerations into account. Future simulation studies might consider how these social determinants intersect with childcare and school configurations. Additionally, this study was primarily concerned with SARS-CoV-2 infection and student days lost. That being said, there are many important outcomes to consider in relation to children's developmental health in the pandemic. Longitudinal studies considering children's learning and mental health outcomes in relation to new childcare and school configurations are strongly indicated [547].

Chapter 6

Conclusion

In the introduction, we stated our intention to “*identify a class of dependable early warning signals [EWS] for coupled disease–infection models of paediatric diseases*”, further specifying our desire to classify, evaluate and compare these EWS. Through computational modelling and statistical analysis, three of the studies detailed in this thesis have established two broad classes of spatial EWS that can be used to predict transitions in the dynamics of both small and large coupled disease-behaviour network models. We established that increases in spatial coherence and aggregation on networks of various structures was indicative of both pending vaccine crises and widespread disease outbreaks in this family of models. We also saw that the interpretation of these warning signals depends heavily on underlying characteristics of the model (especially the function of the included social norm), but that objective EWS performance is robust to noisy dynamics, change in network size and structure, decoupling of dynamics and simple random sampling.

Initial survey

Our first task was to curate a set of established early warning signals for consideration. Despite its heavy basis in graph and network theory, the interdisciplinarity of this work is first seen in the varied origins of the EWS presented: mutual information (Chs. 2, 4) is a tool featured in ecology, finance and statistical physics, and Moran’s I and Geary’s C (Chs. 2, 4) were developed for use with geographical information systems. Join count statistics are used in many fields for correlogram-like point pattern analysis of categorical data but, to our best knowledge, saw little (if any) previous application to the analysis of discrete opinion models. Community detection represents a vast literature in both the fields of social network analysis and graph theory, while percolation theory describes critical phenomena in graph connectivity; from these areas we get

the important notions of community and group formation (Chs 3, 4). Given that our chosen models (and our research questions) sit in the intersection of all these fields, this wide range of adaptable tools and available theory strongly bolstered our confidence in the ultimate utility of a multi-faceted approach.

First study

Reflection on the novelty research question prompted us to develop a simplistic model of the paired dynamics; our parametrised coupling of a $SIRV_p$ (paediatric disease) compartmental model with binary opinion dynamics yielded the organic vaccine crises and disease outbreaks we required, yet was simple enough to allow for the establishment of baseline results for the EWS (facilitating comparison with later studies). Some previous studies remarked on the lack of a definition of epidemics applicable to computer simulation (and broader disease modelling); after literature review, we formulated such a definition, as well as simulation convergence criteria. Though minimal compared to the scope of this thesis, these too represent unique contributions to the literature of our field. Through analysing generated data sets in this first study, we confirmed that noticeable changes in opinion clustering and the synchrony of model dynamics accompanied critical transitions in both model dynamics. Through varying some parameters of the model, we also confirm that these warnings are eventually robust to noise (in the form of random opinion changes), network size and initial aggregate opinion.

Another question we set ourselves initially surrounds comparisons between EWS. As was previously remarked, our desire is to characterise warning behaviour objectively rather than depend on eyesight; this prompted trend analysis and our use of change point detection tests, usually used in the analysis of climatological time series. We feature the results of the Lanzante and the standard normal homogeneity (SNHT) tests throughout the thesis, with two other tests (the Pettitt and Buishand range tests) used strictly for comparison. Application of the two primary change point tests to data series taken at network equilibrium reveal (in many instances) well-defined locations in EWS data series where changes in trend occur. Since the series being tested in all cases were taken from a succession of instances at equilibrium (rather than time series usually used in change point analysis), we interpret the change points yielded by these tests as a confirmation of warnings given by the EWS, as well as a way to quantify the gross movement of the model towards transition. Our use of these change point tests also verified that monitoring the candidate EWS is, in most cases, a better strategy than naïvely monitoring the state of the network. This finding's independence of the strength of the model's social norm (or indeed any other prominent dynamic model characteristic) prompted our omission of this particular comparison from further studies.

Second Study

With this augmented methodology, our second study confirms that changes in the community structure of the simulated social network occurred ‘sufficiently far’ before the critical transitions of the model to yield dependable warning signals of transition. Also, the list of candidate EWS was expanded by measuring the fluctuation of aggregate sentiment, the health of one’s neighbours and the Watts-Strogatz clustering coefficient of similar-opinion networks; these were also confirmed as viable EWS. These results were then taken to comprise a stable basis for comparison, since the third study tests robustness of the EWS to targeted changes in the model. At this stage, we made the decision to discard three candidate EWS used in the work so far: both Moran’s I and Geary’s C (Ch. 2) were confirmed to be linear combinations of join count statistics, and so their further inclusion was seen as duplication. Geary’s C (Ch. 3) and the Getis-Ord G coefficient (Ch. 3, not shown) were discarded due to undependable performance.

Third Study

In the third study, we tested the resilience of the EWS by introducing vaccine hesitance (seen as a delay in the social dynamics) as a separate opinion unenforced by social norm. We also changed the underlying graph structure of the social network from a random graph to that of a network with small world properties, and retrieved the warning signals from random samples of the network (of various sizes). These were seen as ‘real life’ tests for the EWS, similar to those outlined in the social network analysis literature. We also added opinion network diameter and the number of triads on the network as new candidate EWS, and limited change point analysis to using only the SNHT. On the entire network, most EWS performed as well as expected, including the join count and mutual information statistics. As quantified by lead distance, there was a slight decrease in performance when the network was sampled for all EWS.

In sum, we’ve found that mutual information, join count statistics, network diameter, triad census, global clustering coefficient, Moran’s I, the sizes and numbers of opinion communities and echo chambers, the number of changes of opinion, opinion network modularity score and the probability of having an infected social contact all give strong warning signals of impending critical transitions in coupled disease-behaviour models such as ours. Of the previous mentioned EWS, mutual information, join count statistics and the numbers of instances of opinion stand out as the strongest, though ultimately their perceived performance depends on the change point test used to gather their respective lead distances (as can be seen in all three studies in the thesis, different change point tests will net different results).

As for failure of the EWS, all three studies showed occasional failure of all the tools; the locations of these failures (undefined change points in the EWS data series, or a negative lead distance) of

course depended on the change point test used. Perfection was in no way expected; rather we concerned ourselves with trying to identify *patterns* of failure with respect to the social norm. There was no immediately detectable consistent relationship between EWS failure and the strength of the social norm present in any instance; then, the intentional simplicity of our model allows us to dismiss the possibility. Some EWS perform well and others don't, and our results do not allow us to draw any further conclusions on failure of the EWS on this class of simple model; we feel this is best left to a targeted study conducted with a more realistic model.

The results of the third study give us reasonable confidence as to the general strength of these EWS, but establishing suitability for general use requires further study. Previously, we spoke to the simplicity of the model (by design). For this thesis, we weighed two methodologies: use a simple model and obtain incremental results that can be built on, or follow literature with more complex models and obtain results that, while still useful and interesting of their own right, are somewhat isolated. Clearly we thought the first approach would make for a better contribution to the literature of our topic, and therefore carried that through. However, for directions of future work in establishing these EWS, we look to current literature in the broader sphere of empirical social network analysis.

Fourth Study

Our study on the infection dynamics of [COVID-19](#) produced concrete suggestions for mitigation of infection spread, striking a different tone to the other studies' focus on early warnings of critical transition. Through the use of a comparatively simpler ABM, we found that in the cases of both low and high transmission rates, the reduction of class size was (in every scenario) vital to stemming the spread of the disease to the entire network, and especially within the school. More importantly, we can cast this fundamental result in the light of our unifying theme: spatial correlation and connection density facilitates the spread of infection and is therefore useful in predicting spread. This study presented a case where overt manipulation of the connectivity structure was possible, and served to tangentially reinforce the major findings of this thesis not from the point-of-view of early warning signals, but through the underlying principles of location and connectivity.

Limitations and Future Directions

The two foremost restrictive properties of our network are that neighbourhoods are static and identical across dynamics; that is, an individual has the same social (media) contacts as physical contacts, both of which are immutable. This is clearly unrealistic in both ways; the author himself has even done collaborative musical performances with individuals in other countries who

he has never met in person, and encountered many people on public transit who he will never speak to (even before the imposition of these current physical distancing guidelines necessitated by the COVID-19 pandemic). Our physical and social neighbourhoods are not identical, and the loosening of this restriction is the author’s primary concern for future studies. Also, communicative neighbourhoods are rarely static, and so an evolving network (featuring variable rates of communication and dynamic link formation predicated on opinion dynamics) must be considered; humanity’s huge online presence presents the perfect opportunity for monitoring patterns of information exchange and aggregation in large networks, as shown in [129]. The first study in this thesis included simulations run on a network of size $N = 750^2$ (called model V3); the success of the EWS used with this version of the model leads us to believe that the EWS will be insensitive to scale, though we expect computational expense will need to be accounted for.

Other essential tests of the application and deployment of the EWS (and their subsequent use in policy design) would be their performance in ‘real time’; this is the true purpose of change point detection tests, whereas our previous uses were simply confirmatory (with populations at equilibrium). We propose that future studies not only confirm that these EWS give explicit temporal warnings (through some optimal application of change-point tests at each time step), but also the effect of intervention when such a warning signal is received. This is the temporal analogue of our previous (equilibrium) conception of *lead distance*, and gives new meaning to the previously posed question “*how much time do we have to react?*”

The effects of interventions themselves (quarantine, mandatory vaccination, etc) can be studied with such an extended model, but that is outside of the scope of this work. As regards that proposal, observation and evaluation of the EWS investigated in this thesis does not alter the dynamics of the model in any way (as currently conceived), so we cannot reasonably expect the efficacy of any given intervention to be tied to either the EWS monitored or the change point test used.

Summation

In epidemiology, a combination of disease emergence and resurgence and anti-vaccine sentiment proliferation has led to widespread distrust of public health authorities in many countries and far-reaching resistance to healthcare best practices. The boundless efforts of a great many scientists, medical practitioners and politicians has resulted in a large body of work attempting to catalogue the numerous motivations and drivers of vaccine resistance and epidemic events. Unfortunately, the unique ‘perfect storm’ of circumstances exemplified by the current COVID-19 epidemic has exposed multiple disastrous weaknesses in many areas of pandemic preparedness and vaccination policy¹.

¹indeed, *if* and *when* a vaccine is brought to public availability; anti-lockdown protests gathered up anti-vaccine messaging as soon as medical trials of vaccines were announced [522, 324, 507, 32, 72, 217, 421].

Mitigative approaches requiring public cooperation need careful investigation and testing before implementation [423], else an ensuing decline in public trust will lead to non-compliance and resistance possibly quelled only by the destruction wrought by the event itself. An example is the suggested widespread use of (non-medical) facial coverings in the current COVID-19 pandemic; ‘expert’ opinion and public health directive varies across the board due to differing interpretations of context and evidence [78, 261, 293, 189, 345, 246]. This, coupled with complete retraction and reversals of official statements made only days or weeks previously can lead to public outcry, frustration and distrust of national public health systems and officers[372, 448]; some experts assert that the driving force behind adoption and the creation of a strong social norm was ultimately fear [11, 563, 228], which can’t be consistently planned for. This is only a contemporary repeat of situations we have faced in the past, including numerous vaccine scares [399, 142, 321, 241, 242, 133, 236].

The intended use of early warning signals is to dependably (and consistently) inform the enactment of mitigative policies. As outlined above, everyone concerned needs to have strong confidence in the timing and extent of measures. While basic, the results we present in this thesis mark an important contribution to the literature on early warning signals for complex adaptive systems.

References

- [1] Jarle Aarstad, Håvard Ness, and Sven A Haugland. In what ways are small-world and scale-free networks interrelated? In *2013 IEEE International Conference on Industrial Technology (ICIT)*, pages 1483–1487. IEEE, 2013.
- [2] Shaheen A Abdulkareem, Ellen-Wien Augustijn, Yaseen T Mustafa, and Tatiana Filatova. Intelligent judgements over health risks in a spatial agent-based model. *International journal of health geographics*, 17(1):8, 2018.
- [3] Dimitris Achlioptas, Raissa M D’Souza, and Joel Spencer. Explosive percolation in random networks. *Science*, 323(5920):1453–1455, 2009.
- [4] Lada Adamic and Eytan Adar. How to search a social network. *Social Networks*, 27(3):187–203, 2005.
- [5] Lada Adamic, Orkut Buyukkokten, and Eytan Adar. A social network caught in the web. *First monday*, 8(6), 2003.
- [6] Lada A Adamic. The small world web. In *International Conference on Theory and Practice of Digital Libraries*, pages 443–452. Springer, 1999.
- [7] Lada A. Adamic, Thomas M. Lento, Eytan Adar, and Pauline C. Ng. Proceedings of the ninth acm international conference on web search and data mining. In *Information evolution in social networks*, volume 9, pages 473–482. ACM New York, 2016.
- [8] Canada Revenue Agency. Benefits, credits and financial support: CRA and COVID-19, May 2020.
- [9] Samad Mohammad Aghdam and Nima Jaafari Navimipour. Opinion leaders selection in the social networks based on trust relationships proopagation. *Karbala International Journal of Modern Science*, 2:88–97, 2016.

- [10] Syed Faraz Ahmed, Ahmed A Quadeer, and Matthew R McKay. Preliminary identification of potential vaccine targets for the COVID-19 coronavirus (sars-cov-2) based on sars-cov immunological studies. *Viruses*, 12(3):254, 2020.
- [11] Sadaf Ahsan. Why face masks will become the new normal in canada, May 2020. Accessed: 16 May 2020.
- [12] Ali Alessa and Miad Faezipour. Preliminary flu outbreak prediction using twitter posts classification and linear regression with historical centers for disease control and prevention reports: Prediction framework study. *JMIR Public Health and Surveillance*, 5(2):e12383, 2019.
- [13] Murray E Alexander and Seyed M Moghadas. Bifurcation analysis of an sirs epidemic model with generalized incidence. *SIAM Journal on Applied Mathematics*, 65(5):1794–1816, 2005.
- [14] Hans Alexandersson. A homogeneity test applied to precipitation data. *Journal of Climatology*, 6, 1986.
- [15] Hans Alexandersson and Anders Moberg. Homogenization of swedish temperature data. part i: Homogeneity test for linear trends. *International Journal of Climatology*, 17(1), 1986.
- [16] Ahmed Allam, Peter Johannes Schulz, and Kent Nakamoto. The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: two experiments manipulating google output. *Journal of medical internet research*, 16(4):e100, 2014.
- [17] Lucila G. Alvarez-Zuzek, Cristian E. La Rocca, José R. Iglesias, and Lidia A. Braunstein. Epidemic spreading in multiplex networks influenced by opinion exchanges on vaccination. *PLOS ONE*, pages 1–14, 11 2017.
- [18] Luis A Nunes Amaral, Antonio Scala, Marc Barthelemy, and H Eugene Stanley. Classes of small-world networks. *Proceedings of the national academy of sciences*, 97(21):11149–11152, 2000.
- [19] Kristian G Andersen, Andrew Rambaut, W Ian Lipkin, Edward C Holmes, and Robert F Garry. The proximal origin of sars-cov-2. *Nature medicine*, 26(4):450–452, 2020.
- [20] Roy M Anderson, B Anderson, and Robert M May. *Infectious diseases of humans: dynamics and control*. Oxford university press, 1992.
- [21] Roy M. Anderson and Robert M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1992.
- [22] Luc Anselin and Xun Li. Operational local join count statistics for cluster detection. *Journal of geographical systems*, 21(2):189–210, 2019.

- [23] Hélène Arduin, Matthieu Domenech de Cellès, Didier Guillemot, Laurence Watier, and Lulla Opatowski. An agent-based model simulation of influenza interactions at the host level: insight into the influenza-related burden of pneumococcal infections. *BMC Infections Diseases*, 17(382), 2017.
- [24] JR Artalejo and MJ Lopez-Herrero. Stochastic epidemic models: New behavioral indicators of the disease spreading. *Applied Mathematical Modelling*, 38(17-18):4371–4387, 2014.
- [25] Hiroshi Ashikaga and Ameneh Asgari-Targhi. Locating order-disorder phase transition in a cardiac system. *Scientific Reports*, 8(1967), 2018.
- [26] Nicolas Bacaër. *Daniel Bernoulli, d’Alembert and the inoculation of smallpox (1760)*, chapter 4, pages 21–30. Springer, 2011.
- [27] Lars Backstrom and Jure Leskovec. Supervised random walks: predicting and recommending links in social networks. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 635–644, 2011.
- [28] Raul Bagni, Roberto Berchi, and Pasquale Cariello. A comparison of simulation models applied to epidemics. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- [29] Michael T. J. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Charles Griffin & Company Ltd, 2 edition, 1975. "a: chp. 2, sec. 1, pp. 9-10; b: chp. 20, secs. 2-4, pp: 359-366".
- [30] S Balamuralitharan and M Radha. Bifurcation analysis in sir epidemic model with treatment. In *IOP Conf. Series: Journal of Physics: Conf. Series*, volume 1000, 2018.
- [31] Sulzhan Bali, Kearsley A Stewart, and Muhammad Ali Pate. Long shadow of fear in an epidemic: fearonomic effects of ebola on the private sector in nigeria. *BMJ global health*, 1(3):e000111, 2016.
- [32] Philip Ball. Anti-vaccine movement could undermine efforts to end coronavirus pandemic, researchers warn, May 2020. Accessed: 18 May 2020.
- [33] Michele Ballerini, Nicola Cabibbo, Raphael Candelier, Andrea Cavagna, Evaristo Cisbani, Irene Giardina, Vivien Lecomte, Alberto Orlandi, Giorgio Parisi, Andrea Procaccini, et al. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the national academy of sciences*, 105(4):1232–1237, 2008.
- [34] Zoe Bambery, Cynthia H. Cassell, Rebecca E. Bunnell, Kakoli Roy, Zara Ahmed, Rebecca L. Payne, and Martin I. Meltzer. Impact of a hypothetical infectious disease outbreak on us exports and export-based jobs. *Health Security*, 16(1):1–7, 2 2018.

- [35] World Bank. Immunization, dpt (% of children ages 12-23 months) - france. Technical report, The World Bank Group, 2019. Online, retrieved 10-October-2019.
- [36] World Bank. Immunization, hep3 (% of one-year-old children) - france. Technical report, The World Bank Group, 2019. [Online, retrieved 10-October-2019].
- [37] World Bank. Immunization, measles (% of children ages 12-23 months) - france. Technical report, The World Bank Group, 2019. [Online, retrieved 10-October-2019].
- [38] Ziv Bar-Yossef and Maxim Gurevich. Random sampling from a search engine’s index. *Journal of the ACM (JACM)*, 55(5):1–74, 2008.
- [39] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [40] Pablo Barberá, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science*, 26(10):1531–1542, 2015.
- [41] M.S. Bartlett. Measles periodicity and community size. *Journal of the Royal Statistical Society: Series A*, 120(1):48–70, 1957.
- [42] Chris T. Bauch. Imitation dynamics predict vaccinating behaviour. *Proceedings of the Royal Society B*, 272(1573):1669–1675, 2005.
- [43] Chris T. Bauch and Samit Bhattacharyya. Evolutionary game theory and social learning can determine how vaccine scares unfold. *PLOS Computational Biology*, 8(4):1–12, 04 2012.
- [44] Chris T Bauch and Alison P Galvani. Social factors in epidemiology. *Science*, 342(6154):47–49, 2013.
- [45] Chris T Bauch, James O Lloyd-Smith, Megan P Coffee, and Alison P Galvani. Dynamically modeling sars and other newly emerging respiratory illnesses: past, present, and future. *Epidemiology*, pages 791–801, 2005.
- [46] Chris T. Bauch, Ram Sigdel, Joe Pharaon, and Madhur Anand. Early warning signals of regime shifts in coupled human–environment systems. *Proceedings of the National Academy of Sciences of the United States of America*, 113(51):14560–7, 12 2016.
- [47] David Baud, Xiaolong Qi, Karin Nielsen-Saines, Didier Musso, Léo Pomar, and Guillaume Favre. Real estimates of mortality following COVID-19 infection. *The Lancet infectious diseases*, 2020.

- [48] Bert Baumgaertner, Rebecca Tyson, and Stephen Krone. Opinion strength influences the spatial dynamics of opinion formation. *Journal of Mathematical Sociology*, 40(4):207–218, 2016.
- [49] Atiyeh Bayani, Fatemeh Hadaeghi, Sajad Jafari, and Greg Murray. Critical slowing down as an early warning of transitions in episodes of bipolar disorder: A simulation study based on a computational model of circadian activity rhythms. *Chronobiology International*, 34(2):235–245, 2017. PMID: 28060532.
- [50] Jude Bayham and Eli P Fenichel. Impact of school closures for covid-19 on the us health-care workforce and net mortality: a modelling study. *The Lancet Public Health*, 2020.
- [51] Peter S Bearman et al. Chains of affection: The structure of adolescent romantic and sexual networks. In *Handbook of Applied System Science*, pages 180–206. Routledge, 2016.
- [52] Brian Beckage, Louis J. Gross, Katherine Lacasse, Eric Carr, Sara S. Metcalf, Jonathan M. Winter, Peter D. Howe, Nina Fefferman, Travis Franck, Asim Zia, Ann Kinzig, and Forrest M. Hoffman. Linking models of human behaviour and climate alters projected climate change. *Nature Climate Change*, 8:79–84, 01 2016.
- [53] Joshua Becker, Devon Brackbill, and Damon Centola. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences*, 114(26):E5070–E5076, 2017.
- [54] Joshua Becker, Ethan Porter, and Damon Centola. The wisdom of partisan crowds. *Proceedings of the National Academy of Sciences*, 116(22):10717–10722, 2019.
- [55] Niels Becker. Estimation for an epidemic model. *Biometrics*, 32(4):769–777, 12 1976.
- [56] Mariano Beguerisse-Díaz, Guillermo Garduno-Hernández, Borislav Vangelov, Sophia N Yaliraki, and Mauricio Barahona. Interest communities and flow roles in directed networks: the twitter network of the uk riots. *Journal of The Royal Society Interface*, 11(101):20140940, 2014.
- [57] Nefti-Eboni Bempong, Rafael Ruiz de Casteñada, Stefanie Schütte, Isabelle Bolon, Olivia Keiser, Gérard Escher, and Antoine Flahault. Precision global health - the case of ebola: a scoping review. *Journal of Global Health*, 9(1):010404, 6 2019.
- [58] Eshel Ben-Jacob, Inon Cohen, and David L Gutnick. Cooperative organization of bacterial colonies: from genotype to morphotype. *Annual review of microbiology*, 52(1):779–806, 1998.
- [59] IJ Benczik, SZ Benczik, Beate Schmittmann, and RKP Zia. Opinion dynamics on an adaptive random network. *Physical Review E*, 79(4):046104, 2009.

- [60] Daniel Bernoulli. Essai d’une nouvelle analyse de la mortalité causée par la petite vérole, et des avantages de l’inoculation pour la prévenir. *Histoire de l’Acad., Roy. Sci.(Paris) avec Mem*, pages 1–45, 1760. ”Reprinted in: L.P. Bouckaert, B.L. van der Waerden (Eds), Die Werke von Daniel Bernoulli, *Bd. 2 Analysis und Wahrscheinlichkeitsrechnung*, Birkhäuser, Basel, 1982, pp 235. English translation entitled ‘An attempt at a new analysis of the mortality caused by smallpox, and of the advantages of inoculation to prevent it’, in L. Bradley, *Smallpox Inoculation: An Eighteenth Century Mathematical Controversy*, Adult education Department, Nottingham, 1971, p. 21. Reprinted in: S. Haberman, T.A. Sibbett (Eds.) *History of Actuarial Science, vol. VIII, Multiple Decrement and Multiple State Models*, William Pickering, London, 1995, p. 1”.
- [61] Alessandro Bessi, Mauro Coletto, George Alexandru Davidescu, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. Science vs conspiracy: Collective narratives in the age of misinformation. *PloS one*, 10(2):e0118093, 2015.
- [62] Alessandro Bessi, Fabio Petroni, Michela Del Vicario, Fabiana Zollo, Aris Anagnostopoulos, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. Homophily and polarization in the age of misinformation. *The European Physical Journal Special Topics*, 225(10):2047–2059, 2016.
- [63] Cornelia Betsch, Noel T Brewer, Pauline Brocard, Patrick Davies, Wolfgang Gaissmaier, Niels Haase, Julie Leask, Frank Renkewitz, Britta Renner, Valerie F Reyna, et al. Opportunities and challenges of web 2.0 for vaccination decisions. *Vaccine*, 30(25):3727–3733, 2012.
- [64] Cornelia Betsch, Frank Renkewitz, Tilmann Betsch, and Corina Ulshöfer. The influence of vaccine-critical websites on perceiving vaccination risks. *Journal of Health Psychology*, 15(3):446–455, 2010.
- [65] Cornelia Betsch, Frank Renkewitz, and Niels Haase. Effect of narrative reports about vaccine adverse events and bias-awareness disclaimers on vaccine decisions: A simulation of an online patient social network. *Medical Decision Making*, 33(1):14–25, 2013.
- [66] Prantik Bhattacharyya, Ankush Garg, and Shyhtsun Felix Wu. Analysis of user keyword similarity in online social networks. *Social network analysis and mining*, 1(3):143–158, 2011.
- [67] K. Binder. Finite size effects on phase transitions. *Ferroelectrics*, 73:43–67, 1987.
- [68] K. Binder and D.P. Landau. Finite-size scaling at first-order phase transition. *Physical Review B*, 30(3):1477–1485, 8 1984.
- [69] Andrew J Black and Alan J McKane. Stochastic formulation of ecological models and their applications. *Trends in ecology & evolution*, 27(6):337–345, 2012.

- [70] Reinhard Blutner. Dempster-Shafer theory. From the 2011/12 course “Reasoning with Uncertainty”, 2011.
- [71] S Boccaletti, M Ivanchenko, V Latora, A Pluchino, and A Rapisarda. Detecting complex network modularity by dynamical clustering. *Physical Review E*, 75(4):045102, 2007.
- [72] Nicholas Bogel-Burroughs. Antivaccination activists are growing force at virus protests, May 2020. Accessed: 18 May 2020.
- [73] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy USA*, 99(Suppl 3):7280–7287, 5 2002.
- [74] Terry Bossomaier, Lionel Barnett, and Michael Harré. Information and phase transitions in socio-economic systems. *Complex Adaptive Systems Modeling*, 1(1):9, 2013.
- [75] Mondher Bouden, Bernard Moulin, and Pierre Gosselin. The geosimulation of west nile virus propagation: a multi-agent and climate sensitive tool for risk management in public health. *International Journal of Health Geographics*, 7(1):35, 2008.
- [76] Hacène Boukari, Matthew E. Briggs, J. N. Shaumeyer, and Robert W. Gammon. Critical speeding up observed. *Phys. Rev. Lett.*, 65:2654–2657, Nov 1990.
- [77] Andrei Boutyline and Robb Willer. The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3):551–569, 2017.
- [78] Julii Suzanne Brainard, Natalia Jones, Iain Lake, Lee Hooper, and Paul Hunter. Facemasks and similar barriers to prevent respiratory illness such as COVID-19: A rapid systematic review. *medRxiv*, 2020.
- [79] Valentino Braitenberg and Almut Schüz. *Anatomy of the cortex: statistics and geometry*, volume 18. Springer Science & Business Media, 2013.
- [80] Tobias Brett, Marco Ajelli, Quan-Hui Liu, Mary G Krauland, John J Grefenstette, Willem G van Panhuis, Alessandro Vespignani, John M Drake, and Pejman Rohani. Detecting critical slowing down in high-dimensional epidemiological systems. *PLoS computational biology*, 16(3):e1007679, 2020.
- [81] Tobias S. Brett, John M. Drake, and Pejman Rohani. Anticipating the emergence of infectious diseases. *Journal of the Royal Society Interface*, 14(132):20170115, 7 2017.
- [82] Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. Graph structure in the web. *Computer networks*, 33(1-6):309–320, 2000.

- [83] David A. Broniatowski, Michael J. Paul, and Mark Dredze. National and local influenza surveillance through twitter: An analysis of the 2012-2013 influenza epidemic. *PLOS ONE*, 8(12), 12 2013.
- [84] Jo Brown, Amanda J Broderick, and Nick Lee. Word of mouth communication within online communities: Conceptualizing the online social network. *Journal of interactive marketing*, 21(3):2–20, 2007.
- [85] John S Brownstein, Clark C Freifeld, and Lawrence C Madoff. Digital disease detection—harnessing the web for public health surveillance. *New England Journal of Medicine*, 360(21):2153–2157, 2009.
- [86] Axel Bruns. Echo chamber? what echo chamber? reviewing the evidence, 2017. Accessed 17 February 2020.
- [87] Axel Bruns. Filter bubbles and echo chambers: Debunking the myths, Jul 2020.
- [88] Emily K. Brunson. The impact of social networks on parents’ vaccination decisions. *pediatrics*, 131(5), 2013.
- [89] David L Buckeridge, Christian Jauvin, Anya Okhmatovskaia, and Aman D Verma. Simulation analysis platform (snap): a tool for evaluation of public health surveillance and disease control strategies. In *AMIA Annual Symposium Proceedings*, volume 2011, page 161. American Medical Informatics Association, 2011.
- [90] T. A. Buishand. Tests for detecting a shift in the mean of hydrological time series. *Journal of Hydrology*, 73:51–69, 1984.
- [91] Bruno Buonomo. A note on the direction of the transcritical bifurcation in epidemic models. *Nonlinear Anal Model Control*, 20:38–55, 2015.
- [92] Emily A. Burger, Nicole G. Campos, Stephen Sy, Catherine Regan, and Jane J. Kim. Health and economic benefits of single-dose hpv vaccination in a gavi-eligible country. *Vaccine*, 36(32):4823–4829, 2018.
- [93] Alison M Bутtenheim, Sarah T Cherng, and David A Asch. Provider dismissal policies and clustering of vaccine-hesitant families: an agent-based modeling approach. *Human vaccines & immunotherapeutics*, 9(8):1819–1824, 2013.
- [94] Barbara Caci, Maurizio Cardaci, and Marco E Tabacchi. Facebook as a small world: a topological hypothesis. *Social Network Analysis and Mining*, 2(2):163–167, 2012.
- [95] John T Cacioppo, James H Fowler, and Nicholas A Christakis. Alone in the crowd: the structure and spread of loneliness in a large social network. *Journal of personality and social psychology*, 97(6):977, 2009.

- [96] Statistics Canada. Statistics canada 2016 census. <https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/prof/details/page.cfm>.
- [97] Statistics Canada. Table 39-10-0007-01 life expectancy and other elements of the life table, canada and provinces, 2019.
- [98] Kathleen M Carley, Douglas B Fridsma, Elizabeth Casman, Alex Yahja, Neal Altman, Li-Chiou Chen, Boris Kaminsky, and Démian Nave. Biowar: scalable agent-based model of bioattacks. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 36(2):252–265, 2006.
- [99] Dorwin Cartwright and Frank Harary. Structural balance: a generalization of heider’s theory. *Psychological review*, 63(5):277, 1956.
- [100] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Review of Modern Physics*, 81:591–646, 5 2009.
- [101] CDC. Vaccine glossary of terms, Jul 2020.
- [102] Damon Centola, Víctor M Eguíluz, and Michael W Macy. Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications*, 374(1):449–456, 2007.
- [103] Damon Centola and Michael Macy. Complex contagions and the weakness of long ties. *American journal of Sociology*, 113(3):702–734, 2007.
- [104] Edmund Chattoe-Brown. Using agent based modelling to integrate data on attitude change. *Sociological Research Online*, 19(1):159–174, 2014.
- [105] Cecilia Chen. Rebellion against the polio vaccine in nigeria: implications for humanitarian policy. *African Health Sciences*, 4(3):205–208, 2004.
- [106] Hongxu Chen, Hongzhi Yin, Xue Li, Meng Wang, Weitong Chen, and Tong Chen. People opinion topic model: opinion based user clustering in social networks. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1353–1359. International World Wide Web Conferences Steering Committee, 2017.
- [107] Li-Chiou Chen, Boris Kaminsky, Tiffany Tummino, Kathleen M Carley, Elizabeth Casman, Douglas Fridsma, and Alex Yahja. Aligning simulation models of smallpox outbreaks. In *International Conference on Intelligence and Security Informatics*, pages 1–16. Springer, 2004.
- [108] Lijuan Chen, Shouying Huang, Fengde Chen, and Mingjian Fu. The bifurcation analysis and optimal feedback mechanism for an sis epidemic model on networks. *Advances in Difference Equations*, 2019(1):529, 2019.

- [109] Robert T. Chen and Frank DeStefano. Vaccine adverse events: causal or coincidental? *Lancet*, 351(9103):611–612, 1998.
- [110] LiPing Chi. Binary opinion dynamics with noise on random networks. *Chinese science bulletin*, 56(34):3630–3632, 2011.
- [111] Lesley Chiou and Catherine Tucker. Fake news and advertising on social media: A study of the anti-vaccination movement. Technical report, National Bureau of Economic Research, 2018.
- [112] Nakul Chitnis, James M Hyman, and Jim M Cushing. Determining important parameters in the spread of malaria through the sensitivity analysis of a mathematical model. *Bulletin of mathematical biology*, 70(5):1272, 2008.
- [113] Gerardo Chowell, Cécile Viboud, Lone Simonsen, and Seyed M. Moghadas. Characterizing the reproduction number of epidemics with early subexponential growth dynamics. *Journal of the Royal Society Interface*, 13(123):20160659, 10 2016.
- [114] Dimitri A Christakis. School reopening—the pandemic issue that is not getting its due. *JAMA pediatrics*, 2020.
- [115] Nicholas A Christakis and James H Fowler. The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4):370–379, 2007.
- [116] Nicholas A Christakis and James H Fowler. The collective dynamics of smoking in a large social network. *New England journal of medicine*, 358(21):2249–2258, 2008.
- [117] Stéphan Cléménçon, Hector De Arazoza, Fabrice Rossi, and Viet-Chi Tran. Visual mining of epidemic networks. In *International Work-Conference on Artificial Neural Networks*, pages 276–283. Springer, 2011.
- [118] Stéphan Cléménçon, Hector De Arazoza, Fabrice Rossi, and Viet Chi Tran. A statistical network analysis of the hiv/aids epidemics in cuba. *Social Network Analysis and Mining*, 5(1):58, 2015.
- [119] Bruce W. Clements. *Bioterrorism*, chapter 2, pages 27–63. Elsevier, 2009.
- [120] A.D. Cliff, P. Haggett, and J.K Ord. *Spatial Aspects of Influenza Epidemics*. Routledge Hegan & Paul, 2 1997.
- [121] Andres Cliff and Peter Haggett. Methods for the Measurement of Epidemic Velocity from Time-Series Data. *International Journal of Epidemiology*, 11(1):82–89, 03 1982.

- [122] Lucie Cluver, Jamie M Lachman, Lorraine Sherr, Inge Wessels, Etienne Krug, Sabine Rakotomalala, Stephen Blight, Susan Hillis, Gretchen Bachmand, Ohad Green, et al. Parenting in a time of covid-19. *The Lancet*, 2020.
- [123] Richard Colbaugh and Kristin Glass. Predictive analysis for social diffusion: The role of network communities. *arXiv preprint arXiv:0912.5242*, 2009.
- [124] Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson. Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication*, 64:317–332, 2014.
- [125] Shannon Collinson, Kamran Khan, and Jane M. Hefferman. The effects of media reports on disease spread and important public health measurements. *PLOS ONE*, 10(11):e0141423, 2015.
- [126] Anne Condon and Richard M Karp. Algorithms for graph partitioning on the planted partition model. *Random Structures & Algorithms*, 18(2):116–140, 2001.
- [127] Michael D Conover, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. Partisan asymmetries in online political activity. *EPJ Data Science*, 1(1):6, 2012.
- [128] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*, 2011.
- [129] Kimberly Corace, Chatura Prematunge, Anne McCarthy, Rama C. Nair, Virginia Roth, Thomas Hayes, Kathryn Suh, Louise Balfour, and Gary Garber. Predicting influenza vaccination uptake among health care workers: What are the key motivators? *American Journal of Infection Control*, 41(8):679–684, 2013.
- [130] Sophie Cousins. Who hedges its bets: the next global pandemic could be disease x. *BMJ*, 361, 2018.
- [131] CDC COVID and Response Team. Severe outcomes among patients with coronavirus disease 2019 (COVID-19)—united states, february 12–march 16, 2020. *MMWR Morb Mortal Wkly Rep*, 69(12):343–346, 2020.
- [132] Lorenzo Coviello, Massimo Franceschetti, García-Hernandez, and Iyad Rahwan. Predicting and containing epidemic risk using on-line friendship networks. *PloS ONE*, 14(5):e0211765, 5 2019.
- [133] NS Crowcroft and Joseph Britto. Whooping cough—a continuing problem: Pertussis has re-emerged in countries with high vaccination coverage and low mortality, 2002.

- [134] Gabor Csardi, Tamas Nepusz, et al. The igraph software package for complex network research. *InterJournal, Complex Systems*, 1695(5):1–9, 2006.
- [135] D Cucinotta and M Vanelli. Who declares COVID-19 a pandemic. *Acta bio-medica: Atenei Parmensis*, 91(1):157–160, 2020.
- [136] Erik Cuevas. An agent-based model to evaluate the COVID-19 transmission risks in facilities. *Computers in Biology and Medicine*, page 103827, 2020.
- [137] Jingan Cui, Yanan Zhang, and Zhilan Feng. Influence of non-homogeneous mixing on final epidemic size in a meta-population model. *Journal of biological dynamics*, 13(sup1):31–46, 2019.
- [138] Vasilis Dakos, Egbert H. van Nes, D’Odorico, and Marten Scheffer. Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology*, 93(2):264–271, 2012.
- [139] Vasilis Dakos, Egbert H. van Nes, Raú Donangelo, Hugo Fort, and Martin Scheffer. Spatial correlation as leading indicator of catastrophic shifts. *Theoretical Ecology*, 3(3):163–174, 2010.
- [140] Leon Danon, Albert Diaz-Guilera, Jordi Duch, and Alex Arenas. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09):P09008, 2005.
- [141] Leon Danon, Ashley P Ford, Thomas House, Chris P Jewell, Matt J Keeling, Gareth O Roberts, Joshua V Ross, and Matthew C Vernon. Networks and the epidemiology of infectious disease. *Interdisciplinary perspectives on infectious diseases*, 2011, 2011.
- [142] Antonio L Dans, Leonila F Dans, Maria Asuncion A Silvestre, Scott B Halstead, and Gordon H Guyatt. Cause and consequence of loss in vaccine trust. *Human Vaccines & Immunotherapeutics*, 15(3):628, 2019.
- [143] Abhimanyu Das, Sreenivas Gollapudi, and Kamesh Munagala. Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining, WSDM ’14*, pages 403–412. ACM, 2014.
- [144] Moupriya Das and Jason R Green. Critical fluctuations and slowing down of chaos. *Nature communications*, 10(1):2155, 2019.
- [145] Dave. Is the echo chamber idea a myth? *Skeptical Science*, Apr 2018.
- [146] Richard Dawkins. *The Selfish Gene*. Oxford University Press, 3 edition, 2006. 0-19-857519-X.

- [147] Abir De, Isabel Valera, Niloy Ganguly, Sourangshu Bhattacharya, and Manuel Gomez Rodriguez. Learning and forecasting opinion dynamics in social networks. In *30th Conference on Neural Information Processing Systems*, page 9, 2016.
- [148] Joao Marcos De Oliveira and Peter A Gloor. Galaxyscope: Finding the “truth of tribes” on social media. In *Collaborative Innovation Networks*, pages 153–164. Springer, 2018.
- [149] Donald L DeAngelis and Volker Grimm. Individual-based models in ecology after four decades. *F1000prime reports*, 6, 2014.
- [150] Brian Deer. How the case against the MMR vaccine was fixed. *British Medical Journal*, 342(7788):77–82, 2011.
- [151] Guillaume Deffuant. Comparing extremism propagation patterns in continuous opinion models. *Journal of Artificial Societies and Social Simulation*, 9(3), 2006.
- [152] Morris H. Degroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 4 1974.
- [153] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. Echo chambers in the age of misinformation. *arXiv preprint arXiv:1509.00189*, 2015.
- [154] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3):554–559, 2016.
- [155] Michela Del Vicario, Gianna Vivaldo, Alessandro Bessi, Fabiana Zollo, Antinio Scala, Guido Caldarelli, and Walter Quattrociocchi. Echo chambers: Emotional contagion and group polarization on facebook. *Scientific Reports*, 6(37825), 2016.
- [156] Frank DeStefano and Robert T. Chen. Negative association between mmr and autism. *The Lancet*, 353(9169):1987–1988, 06 1999.
- [157] Bruno Di Stefano, Henryk Fuks, and Anna T Lawniczak. Object-oriented implementation of ca/lgca modelling applied to the spread of epidemics. In *2000 Canadian Conference on Electrical and Computer Engineering. Conference Proceedings. Navigating to a New Era (Cat. No. 00TH8492)*, volume 1, pages 26–31. IEEE, 2000.
- [158] Paul Diaz, Paul Constantine, Kelsey Kalmbach, Eric Jones, and Stephen Pankavich. A modified seir model for the spread of ebola in western africa and metrics for resource allocation. *Applied Mathematics and Computation*, 324:141–155, 2018.

- [159] Catherine Dibble, Stephen Wendel, and Kristopher Carle. Simulating pandemic influenza risks of us cities. In S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, Tew J. D., and R. R. Barton, editors, *Proceedings of the 2007 Winter Simulation Conference*, pages 1548–1550, 2007.
- [160] Kenne A Dibner, Heidi A Schweingruber, and Dimitri A Christakis. Reopening k-12 schools during the covid-19 pandemic: A report from the national academies of sciences, engineering, and medicine. *JAMA*, 2020.
- [161] Klaus Dietz and J. A. P. Heesterbeek. Daniel Bernoulli’s epidemiological model revisited. *Mathematical Biosciences*, 180:1–21, 2002.
- [162] Cees Diks, Cars Hommes, and Juanxi Wang. Critical slowing down as an early warning signal for financial crises. *Empirical Economics*, 2018.
- [163] Paul Donangelo, Hugo Fort, Vasilis Dakos, Marten Scheffer, and Egbert H. van Nes. Early warnings for catastrophic shifts in ecosystems: Comparison between spatial and temporal indicators. *International Journal of Bifurcation and Chaos*, 20(2):315–321, 2010.
- [164] C. Patrick Doncaster, Vasthi Alonso Chávez, Clément Viguier, Rong Wang, Enlou Zhang, Xuhui Dong, John A. Dearing, Peter G. Langdon, and James G. Dyke. Early warning of critical transitions in biodiversity from compositional disorder. *Ecological Society of America*, 97(11), 8 2016.
- [165] John M. Drake, Tobias S. Brett, Shiyang Chen, Bogdan I. Epureanu, Matthew J. Ferrari, Éric Marty, Paige B. Miller, Eamon B. O’Dea, Suzanne M. O’Regan, Andrew W. Park, and Pejman Rohani. The statistics of epidemic transmission. *PLOS Computational Biology*, 15(5):1–14, 5 2019.
- [166] Eve Dubé, Caroline Laberge, Maryse Guay, Paul Bramadat, Réal Roy, and Julie A. Bettinger. Vaccine hesitancy: an overview. *Human Vaccines & Immunotherapeutics*, 9(8):1763–1773, 8 2013.
- [167] Elizabeth Dubois and Grant Black. The myth of the echo chamber. *The Conversation*, Mar 2018.
- [168] Maeve Duggan and Aaron Smith. *The Political Environment on Social Media: Some Users Enjoy the Opportunities for Political Debate and Engagement that Social Media Facilitates, But Many More Express Resignation, Frustration Over the Tone and Content of Social Platforms*. Pew Research Center, 2016.
- [169] Jill Bigley Dunham. An agent-based spatially explicit epidemiological model in mason. *Journal of Artificial Societies and Social Simulation*, 9(1), 2005.

- [170] Burkhard Dünweg, David P Landau, and Andrey I Milchev. *Computer simulations of surfaces and interfaces*, volume 114. Springer Science & Business Media, 2013.
- [171] William H Dutton, Bianca Reisdorf, Elizabeth Dubois, and Grant Blank. Social shaping of the politics of internet search and networking: Moving beyond filter bubbles, echo chambers, and fake news. *Quello Center Working Paper*, 2017.
- [172] Owen Dyer. COVID-19: Trump stokes protests against social distancing measures, 2020.
- [173] Alfinito. E., M. Beccaria, and G. Macorini. Critical behavior in a stochastic model of vector mediated epidemics. *Scientific Reports*, 6(27202), 6 2016.
- [174] David Easley, Jon Kleinberg, et al. Networks, crowds, and markets: Reasoning about a highly connected world. *Significance*, 9:43–44, 2012.
- [175] David R. Easterling and Thomas C. Paterson. A new method for detecting undocumented discontinuities in climatological time series. *International Journal in Climatology*, 15(4):369–377, 1995.
- [176] Mostafa El-Bermawy. Your filter bubble is ruining democracy. *Wired: Business*, 2016.
- [177] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(1):17–60, 1960.
- [178] Kamil Erguler, Nastassya L Chandra, Yiannis Proestos, Jos Lelieveld, George K Christophides, and Paul E Parham. A large-scale stochastic spatiotemporal model for aedes albopictus-borne chikungunya epidemiology. *PloS one*, 12(3), 2017.
- [179] E Yagmur Erten, Joseph T Lizier, Mahendra Piraveenan, and Mikhail Prokopenko. Criticality and information dynamics in epidemiological models. *Entropy*, 19(5):194, 2017.
- [180] Fatima Etbaigha, Allan R Willms, and Zvonimir Poljak. An seir model of influenza a virus infection and reinfection within a farrow-to-finish swine farm. *PLOS one*, 13(9), 2018.
- [181] Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.
- [182] Gunther Eysenbach. Infodemiology: tracking flu-related searches on the web for syndromic surveillance. In *AMIA Annual Symposium Proceedings*, volume 2006, page 244. American Medical Informatics Association, 2006.
- [183] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. *ACM SIGCOMM computer communication review*, 29(4):251–262, 1999.

- [184] Victoria Y. Fan, Dean T. Jamison, and Lawrence H. Summers. Pandemic risk: how large are the expected losses? *Bulletin of the World Health Organization*, 96:129–134, 2018.
- [185] Zhanpeng Fang and Je Tang. Uncovering the formation of triadic closure on social networks. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [186] Anthony S Fauci, H Clifford Lane, and Robert R Redfield. COVID-19—navigating the uncharted, 2020.
- [187] Scott L Feld. Social structural determinants of similarity among associates. *American Sociological Review*, 47(6):797–801, 1982.
- [188] Dan S Felsenthal and Moshé Machover. Ternary voting games. *International journal of game theory*, 26(3):335–351, 1997.
- [189] Shuo Feng, Chen Shen, Nan Xia, Wei Song, Mengzhen Fan, and Benjamin J Cowling. Rational use of face masks in the COVID-19 pandemic. *The Lancet Respiratory Medicine*, 8(5):434–436, 2020.
- [190] Neil M Ferguson, Derek AT Cummings, Christophe Fraser, James C Cajka, Philip C Cooley, and Donald S Burke. Strategies for mitigating an influenza pandemic. *Nature*, 442(7101):448–452, 2006.
- [191] Emilio Ferrara, Pasquale De Meo, Giacomo Fiumara, and Robert Baumgartner. Web data extraction, applications and techniques: A survey. *Knowledge-based systems*, 70:301–323, 2014.
- [192] Fernando Ferri, Patrizia Grifoni, and Tiziana Guzzo. New forms of social and professional digital relationships: the case of facebook. *Social network analysis and mining*, 2(2):121–137, 2012.
- [193] Graziela P. Figueredo, Peer-Olaf Siebers, Markus R. Owen, Jenna Reys, and Uwe Aickelin. Comparing stochastic differential equations and agent-based modelling and simulation for early-stage cancer. *PLOS ONE*, 9(4):e95150, 04 2014.
- [194] Graziano Fiorillo, Paolo Bocchini, and Javier Buceta. A predictive spatial distribution framework for filovirus-infected bats. *Scientific Reports*, 8(7970), 5 2018.
- [195] M.E. Fischer. The theory of equilibrium critical phenomena. *Reports on Progress in Physics*, 30:615–730, 1967.
- [196] M.M. Fischer. *Spatial Autocorrelation*, chapter B.3, pages 255–300. Springer Science & Business Media, 2009. Equation B.3.7 has a mistake in its numerator: there is a stray minus sign between the $(y_i - \bar{y})$ and $(y_j - \bar{y})$ terms.

- [197] B Fischhoff. Human decision making, in: Research needs for human factors. In *Community of Human Factors*. National Academy Press, 1983.
- [198] Seth Flaxman, Sharad Goel, and Justin M. Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(Special Issue):298–320, 2016.
- [199] Centers for Disease Control and Prevention. Disease surveillance and monitoring. U.S. Department of Health and Human Services, 2018.
- [200] Centres for Disease Control and Prevention (CDC). *Introduction to Epidemiology*, chapter 1, pages 72–80. U.S. Department of Health and Human Services, Atlanta, GA 30333, 2012 edition, 5 2012. a self-study course.
- [201] Marie-Josée Fortin, Mark R. T. Dale, and Jay ver Hoef. *Spatial analysis in ecology*, chapter 4, pages 2051–2058. John Wiley & Sons, Ltd, Chichester, 2002.
- [202] GLobal Community Foundation. Digital media literacy: What is an echo chamber?, 2020.
- [203] James H Fowler and Nicholas A Christakis. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study. *Bmj*, 337:a2338, 2008.
- [204] Christophe Fraser, Steven Riley, Roy M Anderson, and Neil M Ferguson. Factors that make an infectious disease outbreak controllable. *Proceedings of the National Academy of Sciences*, 101(16):6146–6151, 2004.
- [205] L. Freed, Samuel L. Katz, and Sarah J. Clark. Safety of vaccinations: Miss America, the media, and public health. *JAMA*, 276(33):1869–1872, 1996.
- [206] Clark C. Freifeld, Kenneth D. Mandl, Ben Y. Reis, and John S. Brownstein. Healthmap: Global infectious disease monitoring through automated classification and visualization of internet media reports. *Journal of the America Medical Informatics Association*, 15(2):150–157, 4 2008.
- [207] John R.P. French Jr. A formal theory of social power. *Psychological Review*, 63(3):181–194, 1956.
- [208] Enrique Frías-Martinez, Graham Williamson, and Vanessa Frías-Martinez. An agent-based model of epidemic spread using human mobility and social network information. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*, pages 57–64. IEEE, 2011.
- [209] Feng Fu, Daniel I Rosenbloom, Long Wang, and Martin A Nowak. Imitation dynamics of vaccination behaviour on social networks. *Proceedings of the Royal Society B: Biological Sciences*, 278(1702):42–49, 2011.

- [210] Serge Galam. Minority opinion spreading in random geometry. *The European Physical Journal B - Condensed Matter and Complex Systems*, 25(4):403–406, 2 2002.
- [211] Serge Galam. *Sociophysics: A physicist’s modeling of psycho-political phenomena*. Understanding Complex Systems. Springer, New York, 2012.
- [212] E. J. Gangarosa, A. M. Galazka, L. M. Phillips, R. E. Gangarosa, E. Miller, and R. T. Chen. Impact of anti-vaccine movements on pertussis control: the untold story. *Lancet*, 351(9099):356–361, 1998.
- [213] Jason R Gantenberg, Maximilian King, Madeline C Montgomery, Omar Galárraga, Mattia Prosperi, Philip A Chan, and Brandon DL Marshall. Improving the impact of hiv pre-exposure prophylaxis implementation in small urban centers among men who have sex with men: An agent-based modelling study. *PloS one*, 13(7):e0199915, 2018.
- [214] Junling Gao, Pinpin Zheng, Yingnan Jia, Hao Chen, Yimeng Mao, Suhong Chen, Yi Wang, Hua Fu, and Junming Dai. Mental health problems and social media exposure during COVID-19 outbreak. *Plos one*, 15(4):e0231924, 2020.
- [215] Geoffrey P Garnett, Simon Cousens, Timothy B Hallett, Richard Steketee, and Neff Walker. Mathematical models in the evaluation of health programmes. *The Lancet*, 378(9790):515–525, 2011.
- [216] R. Kelly Garrett. Echo chambers online?: Politically motivated selective exposure among internet news users. *Journal of Computer-Mediated Communication*, 14(2):265–285, 2009.
- [217] Mariel Garza. Opinion: Trump is handing anti-vaxxers an invitation to smear coronavirus vaccines, May 2020. Accessed: 18 May 2020.
- [218] Jacques Gautrais, Christian Jost, and Guy Theraulaz. Key behavioural factors in a self-organised fish school model. In *Annales Zoologici Fennici*, volume 45, pages 415–428. BioOne, 2008.
- [219] Bertram Gawronski. Back to the future of dissonance theory: Cognitive consistency as a core motive. *Social Cognition*, 30(6):652–668, 2012.
- [220] R.C. Geary. The contiguity ratio and statistical mapping. *The Incorporated Statistician*, 5(3):”115–127, 129–146”, 11 1954.
- [221] Arthur Getis. A history of the concept of spatial autocorrelation: A geographer’s perspective. *Geographical Analysis*, 40(3), 2008.
- [222] Wayne M. Getz and John Pickering. Epidemic models: Thresholds and population regulation. *The American Naturalist*, 121(6):892–898, 6 1983.

- [223] Nabeel Gillani, Ann Yuan, Martin Saveski, Soroush Vosoughi, and Deb Roy. Me, my echo chamber, and i: introspection on social media polarization. In *Proceedings of the 2018 World Wide Web Conference*, pages 823–831, 2018.
- [224] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the national academy of sciences*, 99(12):7821–7826, 2002.
- [225] Matthew Glomski and Edward Ohanian. Eradicating a disease: Lessons from mathematical epidemiology. *The College Mathematics Journal*, 43(2):123–132, 2012.
- [226] Fiona Godlee, Jane Smith, and Harvey Marcovitch. Wakefield’s article linking mmr vaccine and autism was fraudulent. *British Medical Journal*, 342, 2011.
- [227] Neal D Goldstein, Michael T LeVasseur, Nguyen K Tran, Jonathan Purtle, Seth L Welles, and Stephen C Eppes. Modeling hpv vaccination scale-up among urban young men who have sex with men in the context of hiv. *Vaccine*, 2019.
- [228] Mark Gollom. Why attitudes toward mask wearing are quickly changing in canada and the u.s. — cbc news, May 2020. Accessed: 16 May 2020.
- [229] Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):21:1–21:37, feb 2012.
- [230] Steven M Goodreau, James A Kitts, and Martina Morris. Birds of a feather, or friend of a friend? using exponential random graph models to investigate adolescent social networks. *Demography*, 46(1):103–125, 2009.
- [231] Dimitrios Gouglas, Tung Thanh Le, Klara Henderson, Aristidis Kaloudis, Trygve Danielsen, Nicholas Caspersen Hammersland, James M Robinson, Penny M Heaton, and John-Arne Røttingen. Estimating the cost of vaccine development against epidemic infectious diseases: a cost minimisation study. *The Lancet Global Health*, 6(12):e1386 – e1396, 2018.
- [232] Peter Grabowski. Social media isn’t polarizing us. *Medium*, Aug 2020.
- [233] Amanda Graham, Frank Cullen, Justin Pickett, Cheryl Jonson, Murat Haner, and Melissa Sloan. Faith in trump, moral foundations, and social distancing defiance during the coronavirus pandemic. *Moral Foundations, and Social Distancing Defiance During the Coronavirus Pandemic (April 22, 2020)*, 2020.
- [234] P. Grassberger. On the critical behaviour of the general epidemic process and dynamical percolation. *Mathematical Biosciences*, 63:157–172, 1983.
- [235] David Robert Grimes. Echo chambers are dangerous – we must try to break free of our online bubbles. *The Guardian*, Dec 2017.

- [236] Liza Gross. A broken trust: Lessons from the vaccine–autism wars: Researchers long ago rejected the theory that vaccines cause autism, yet many parents don’t believe them. can scientists bridge the gap between evidence and doubt? *PLoS Biology*, 7(5), 2009.
- [237] Shi-Jian Gu, Chang-Pu Sun, and Hai-Qing Lin. Universal role of correlation entropy in critical phenomena. *Journal of Physics A: Mathematical and Theoretical*, 41(2), 12 2007.
- [238] Fiona M. Guerra, Shelly Bolotin, Gillian Lim, Jane Heffernan, Shelley L. Deeks, Ye Li, and Natasha S. Crowcroft. The basic reproduction number (r_0) of measles: a systematic review. *The Lancet Infectious Diseases*, 17(12):3420–e428, 2017.
- [239] Pedro Calais Guerra, Wagner Meira Jr, Claire Cardie, and Robert Kleinberg. A measure of polarization on social media networks based on community boundaries. In *Seventh International AAAI Conference on Weblogs and Social Media*, 2013.
- [240] Douglas Guilbeault, Joshua Becker, and Damon Centola. Social learning and partisan bias in the interpretation of climate trends. *Proceedings of the National Academy of Sciences*, 115(39):9714–9719, 2018.
- [241] Louise Guillaume. *The measles, mumps and rubella (MMR) vaccine scare: the information needs and information sources of parents*. PhD thesis, University of Sheffield, 2006.
- [242] Louise Guillaume and Peter A Bath. A content analysis of mass media sources in relation to the mmr vaccine scare. *Health Informatics Journal*, 14(4):323–334, 2008.
- [243] Shashank Gupta, Raghuveer Thirukovalluru, Manjira Sinha, and Sandya Mannarswamy. Cimtdetect: a community infused matrix-tensor coupled factorization based method for fake news detection. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 278–281. IEEE, 2018.
- [244] Ryan J Hammer, Thomas W Moore, Patrick D Finley, and Robert J Glass. The role of community structure in opinion cluster formation. In *International Conference on Complex Sciences*, pages 127–139. Springer, 2012.
- [245] Sharon JB Hanley, Eiji Yoshioka, Yoshiya Ito, and Reiko Kishi. HPV vaccination crisis in japan. *The Lancet*, 385(9987):2571, 2015.
- [246] Gail Harding. Use of masks during COVID-19 not always the answer, says epidemiologist — cbc news, May 2020. Accessed: 10 May 2020.
- [247] Nathan Harding, Ramil Nigmatullin, and Mikhail Prokopenko. Thermodynamic efficiency of contagions: a statistical mechanical analysis of the sis epidemic model. *Interface focus*, 8(6):20180036, 2018.

- [248] Irene A. Harmsen, Liesbeth Mollema, Robert A. C. Ruiter, Theo G. W. Paulussen, Hester E. de Melker, and Gerjo Kok. Why parents refuse childhood vaccination: a qualitative study using online focus groups. *BMC Public Health*, 13(1):1183, 2013.
- [249] Daihai He, Edward L. Ionides, and Aaron A. King. Plug-and-play inference for disease dynamics: measles in large and small populations as a case study. *Journal of the Royal Society Interface*, 7:271–283, 2009.
- [250] Laura Heavey, Geraldine Casey, Ciara Kelly, David Kelly, and Geraldine McDarby. No evidence of secondary transmission of covid-19 from children attending school in ireland, 2020. *Eurosurveillance*, 25(21):2000903, 2020.
- [251] Rainer Hegselmann and Ulrich Krause. Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25(4):381–405, 2005.
- [252] Pri Hermawan, Yuliati Komar, Utomo Sarjono Putro, and Soehartati Gondowiardjo. An agent-based simulation approach for government intervention in reducing high-risk hpv. In *Agent-Based Approaches in Economics and Social Complex Systems IX*, pages 27–43. Springer, 2017.
- [253] Jose L. Herrera, Ravi Srinivasan, John S. Brownstein, Alison P. Galvani, and Lauren Ancel Meyers. Disease surveillance on complex networks. *PLOS Computational Biology*, 12(7):e1004928, 7 2016.
- [254] John C. Hershey, David A. Asch, Thi Thumasathit, Jacqueline Meszaros, and Victor V. Waters. The roles of altruism, free riding, and bandwagoning in vaccination decisions. *Organizational Behavior and Human Decision Processes*, 59(2):177–187, 1994.
- [255] Herbert W Hethcote and James A Yorke. A simple model for gonorrhea dynamics. In *Gonorrhea Transmission Dynamics and Control*, pages 18–24. Springer, 1984.
- [256] . Edward C. Holmes, Andrew Rambaut, and Kristian G. Andersen. Pandemics: spend on surveillance, not prediction. *Nature*, 558:180–182, 2018.
- [257] H. Honda, Y. Chimizu, and M. Rutter. No effect of mmr withdrawal on the incidence of autism: a total population study. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, 46(6):572–579, 2005.
- [258] Richard Horton. The lessons of MMR. *Lancet*, 363(2004):747–749, 2004.
- [259] Kartik Hosanagar. Blame the echo chamber on facebook. but blame yourself, too. *Wired*, Nov 2016.
- [260] Peter J Hotez. Covid19 meets the antivaccine movement. *Microbes and Infection*, 2020.

- [261] Jeremy Howard, Austin Huang, Zhiyuan Li, Zeynep Tufekci, Vladimir Zdimal, Helene-Mari van der Westhuizen, Arne von Delft, Amy Price, Lex Fridman, Lei-Han Tang, et al. Face masks against COVID-19: an evidence review, 2020.
- [262] Da-Yin Hua and Ke Gao. Phase transitions of an epidemic spreading model in small-world networks. *Communications in Theoretical Physics*, 55(6):1127–1131, 6 2011.
- [263] Hong Huang, Jie Tang, Lu Liu, JarDeo Luo, and Ziaoming Fu. Triadic closure pattern analysis and prediction in social networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(12), 12 2015.
- [264] Anette Hulth, Gustaf Rydevik, and Annika Linde. Web queries as a source for syndromic surveillance. *PloS one*, 4(2):e4378, 2009.
- [265] Elizabeth Hunter, Brian Mac Namee, and John D Kelleher. A comparison of agent-based models and equation based models for infectious disease epidemiology. In *AICS*, pages 33–44, 2018.
- [266] Hunter-Elizabeth, Brian Mac Namee, and John Kelleher. An open-data-driven agent-based model to simulate infectious disease outbreaks. *PLOS ONE*, 14(1):e0211245, 2018.
- [267] Amit Huppert and G Katriel. Mathematical modelling and prediction in infectious disease epidemiology. *Clinical microbiology and infection*, 19(11):999–1005, 2013.
- [268] D. Isaacs, H. A. Kilham, and H. Marshall. Should routine childhood immunization be compulsory? *Journal of Paediatric Child Health*, 40:392–396, 2004.
- [269] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, New Jersey, 2008.
- [270] R. K. Jaiswal, A. K. Lohani, and H. L. Tiwari. Statistical analysis for change detection and trend assessment in climatological parameters. *Environmental Processes*, 2:729–749, 8 2015.
- [271] Mahdi Jalili and Matjaž Perc. Information cascades in complex networks. *Journal of Complex Networks*, 5:665–693, 7 2017.
- [272] Hans-Karl Janssen and Olaf Stenull. First-order phase transitions in outbreaks of co-infectious diseases and the extended general epidemic process. *EPL (Europhysics Letters)*, 113(2):26005, 2 2016.
- [273] Bertrand Jayles, Hye-rin Kim, Ramón Escobedo, Stéphane Cezera, Adrien Blanchet, Tatsuya Kameda, Clément Sire, and Guy Theraulaz. How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625, 2017.

- [274] Peter C. Jentsch, Madhur Anand, and Chris T. Bauch. Spatial autocorrelation as an early warning signal of regime shifts in a multiplex disease-model behaviour network. *Journal of Theoretical Biology*, 448:17–25, 7 2018.
- [275] Terry C Jones, Barbara Mühlemann, Talitha Veith, Guido Biele, Marta Zuchowski, Jörg Hoffmann, Angela Stein, Anke Edelmann, Victor Max Corman, and Christian Drosten. An analysis of sars-cov-2 viral load by patient age. *medRxiv*, 2020.
- [276] Tracey Elizabeth Claire Jones-Konneh, Tomomi Suda, Hiroyuki Sasaki, and Shinichi Egawa. Agent-based modeling and simulation of nosocomial infection among healthcare workers during ebola virus disease outbreak in sierra leone. *The Tohoku journal of experimental medicine*, 245(4):231–238, 2018.
- [277] Richard Kahn and Douglas Kellner. New media and internet activism: from the ‘battle of seattle’ to blogging. *New media & society*, 6(1):87–95, 2004.
- [278] Nathan Kallus. Predicting crowd behavior with big public data. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 625–630, 2014.
- [279] Yun-Jung Kang. Mortality rate of infection with COVID-19 in korea from the perspective of underlying disease. *Disaster Medicine and Public Health Preparedness*, pages 1–3, 2020.
- [280] Rune Karlsen, Kari Steen-Johnsen, Dag Wollenbæk, and Bernard Enjolras. Echo chamber and trench warfare dynamics in online debates. *European Journal of Communication*, 0(0):0267323117695734, 2017.
- [281] Matt J Keeling and Ken TD Eames. Networks and epidemic models. *Journal of the Royal Society Interface*, 2(4):295–307, 2005.
- [282] Matt J. Keeling and B.T. Grenfell. Disease extinction and community size: Modeling the persistence of measles. *Science*, 275(5296):65–67, 1997.
- [283] Sonia Kefi, Vishweshia Guttal, William A Brock, Stephen R Carpenter, Aaron M Ellison, Valerie N Livina, David A Seekell, Marten Scheffer, Egbert H van Nes, and Vasilis Dakos. Early warning signals of ecological transitions: methods for spatial patterns. *PloS one*, 9(3), 2014.
- [284] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’03, pages 137–146, New York, NY, USA, 2003. ACM.
- [285] Eben Kenah and Joel C. Miller. Epidemic percolation networks, epidemic outcomes, and interventions. *Interdisciplinary Perspectives on Infectious Diseases*, 2011:1–13, 2011.

- [286] Eben Kenah and James M. Ronins. Second look at the spread of epidemics on networks. *Physical Review E*, 76(036113), 2007.
- [287] Ali S. Khan and Nicki Pesik. *Forensic Public Health, Epidemiological and Microbiological Investigations for Biosecurity*, pages 239–256. Elsevier, 2 edition, 2011.
- [288] Mayukh Nilay Khan and Jackie C. May. Phase transitions in random graphs - outbreak of epidemics to network robustness and fragility. hosted by the University of Illinois Urbana-Champaign, accessed 30-05-2019, 5 2010.
- [289] Masahiro Kimura, Kazumi Saito, Kouzou Ohara, and Hiroshi Motoda. Learning to predict opinion share in social networks. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI'10*, pages 1364–1370. AAAI Press, 2010.
- [290] István Z Kiss, Joel C Miller, Péter L Simon, et al. Mathematics of epidemics on networks. *Cham: Springer*, 598, 2017.
- [291] Mikko Kivelä, Alex Arenas, Marc Barthélemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. Multilayer networks. *Journal of Complex Networks*, 2(3):203, 2014.
- [292] David Klepper and Beatrice Dupuy. Conspiracies are spreading over coronavirus vaccines before one even exists, Apr 2020. Accessed: 7 June 2020.
- [293] Michael Klompas, Charles A Morris, Julia Sinclair, Madelyn Pearson, and Erica S Shenoy. Universal masking in hospitals in the COVID-19 era. *New England Journal of Medicine*, 2020. Accessed: 18 May 2020.
- [294] Caprice Knapp, Vanessa Madden, Hua Wang, Phyllis Sloyer, and Elizabeth Shenkman. Internet use and ehealth literacy of low-income parents whose children have special health care needs. *Journal of medical Internet research*, 13(3):e75, 2011.
- [295] David Knoke and Song Yang. *Social Network Analysis*. SAGE Publications, Inc., 2 edition, 2008.
- [296] Wee Chian Koh, Lin Naing, Muhammad Ali Rosledzana, Mohammad Fathi Alikhan, Liling Chaw, Matthew Griffith, Roberta Pastore, and Justin Wong. What do we know about sars-cov-2 transmission? a systematic review and meta-analysis of the secondary attack rate, serial interval, and asymptomatic infection. *medRxiv*, 2020.
- [297] Gueorgi Kossinets and Duncan J Watts. Empirical analysis of an evolving social network. *science*, 311(5757):88–90, 2006.
- [298] P.L. Krapivsky and S. Redner. Dynamics of majority rule in two-state interacting spin systems. *Physical Review Letters*, 90(23):238701, 6 2003.

- [299] Mirjam Kretzschmar, Yvonne THP van Duynhoven, and Anton J Severijnen. Modeling prevention strategies for gonorrhoea and chlamydia using stochastic network simulations. *American Journal of Epidemiology*, 144(3):306–317, 1996.
- [300] P.V. Krishna Iyer. The first and second moments of some probability distributions arising from points on a lattice and their application. *Biometrika*, 36(1/2):135–141, 6 1948.
- [301] Christian Kuehn. A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics. *Physica D: Nonlinear Phenomena*, 240(12):1020 – 1035, 2011.
- [302] Christian Kuehn. A mathematical framework for critical transitions: Normal forms, variance and applications. *Journal of Nonlinear Science*, 23(3):457–510, 6 2013.
- [303] Ravi Kumar, Mohammad Mahdian, and Mary McGlohon. Dynamics of conversations. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 553–562. ACM, 2010.
- [304] Ravi Kumar, Jasmine Novak, and Andrew Tomkins. Structure and evolution of online social networks. In *Link mining: models, algorithms, and applications*, pages 337–357. Springer, 2010.
- [305] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. Trawling the web for emerging cyber-communities. *Computer networks*, 31(11-16):1481–1493, 1999.
- [306] Supriya Kumar, John J Grefenstette, David Galloway, Steven M Albert, and Donald S Burke. Policies to reduce influenza in the workplace: impact assessments using an agent-based model. *American journal of public health*, 103(8):1406–1411, 2013.
- [307] Ziva Kunda. The case for motivated reasoning. *Psychological bulletin*, 108(3):480, 1990.
- [308] Mark Kutner, Elizabeth Greenburg, Ying Jin, and Christine Paulsen. The health literacy of america’s adults: Results from the 2003 national assessment of adult literacy. nces 2006-483. *National Center for Education Statistics*, 2006.
- [309] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600, 2010.
- [310] Alexander Lachmann. Correcting under-reported covid-19 case numbers. *medRxiv*, 2020.
- [311] Jennifer Lackey. True story: Echo chambers are not the problem. *Morning Consult*, Nov 2018.

- [312] Steven J. Lade and Thilo Gross. Early warning signals for critical transitions: A generalized modeling approach. *PLOS Computational Biology*, 8(2):e10002360, 2 2012.
- [313] C. Lagorio, M. Dickison, F. Vazquez, L.A. Braunstein, P.A. Macri, M.V. Migueles, S. Havlin, and H.E. Stanley. Quarantine-generated phase transition in epidemic spreading. *Physical Review E*, 83(026102):1–6, 2 2011.
- [314] Andrea Lancichinetti and Santo Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E*, 80(1):016118, 2009.
- [315] Andrea Lancichinetti and Santo Fortunato. Community detection algorithms: a comparative analysis. *Physical review E*, 80(5):056117, 2009.
- [316] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. Benchmark graphs for testing community detection algorithms. *Physical review E*, 78(4):046110, 2008.
- [317] Andrea Lancichinetti, Mikko Kivelä, Jari Saramäki, and Santo Fortunato. Characterizing the community structure of complex networks. *PloS one*, 5(8):e11976, 2010.
- [318] Sarah Lane, Noni E MacDonald, Melanie Marti, and Laure Dumolard. Vaccine hesitancy around the globe: Analysis of three years of who/unicef joint reporting form data-2015–2017. *Vaccine*, 36(26):3861–3867, 2018.
- [319] John R. Lanzante. Resistant, robust and non-parametric techniques for the analysis of climate data: Theory and examples, including applications to historical radiosonde station data. *International Journal of Climatology*, 61(11), 1996.
- [320] Vincent Laperriere, Dominique Badariotti, Arnaud Banos, and Jean-Pierre Müller. Structural validation of an individual-based model for plague epidemics simulation. *Ecological Complexity*, 6(2):102–112, 2009.
- [321] Heidi J Larson, Kenneth Hartigan-Go, and Alexandre de Figueiredo. Vaccine confidence plummets in the philippines following dengue vaccine scare: why it matters to pandemic preparedness. *Human vaccines & immunotherapeutics*, 15(3):625–627, 2019.
- [322] Marek Laskowski, Bryan C. Demianyuk, Julia Witt, Shamir N. Mukhi, Marcia R. Friesen, and Robert D. McLeod. Agent-based modeling of the spread of influenza-like illness in an emergency department: A simulation study. *IEEE Transactions on Information Technology in Biomedicine*, 15(6):877–889, 2011.
- [323] Lionel Laurent. The COVID-19 vaccine fight is getting ugly, May 2020. Accessed: 7 June 2020.

- [324] Tara Law. COVID-19 vaccine faces challenge of skeptics, misinformation, May 2020. Accessed: 18 May 2020.
- [325] David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. The parable of google flu: Traps in big data analysis. *Science*, 343:1203–1205, 3 2014.
- [326] T Thanh Le, Zacharias Andreadakis, Arun Kumar, R Gomez Roman, Stig Tollefsen, Melanie Saville, and Stephen Mayhew. The COVID-19 vaccine development landscape. *Nat Rev Drug Discov*, 19(5):305–6, 2020.
- [327] Bruce Y Lee, Shawn T Brown, Philip Cooley, Maggie A Potter, William D Wheaton, Ronald E Voorhees, Samuel Stebbins, John J Grefenstette, Shanta M Zimmer, Richard Zimmerman, et al. Simulating school closure strategies to mitigate an influenza epidemic. *Journal of public health management and practice: JPHMP*, 16(3):252, 2010.
- [328] Wookey Lee, James J Lee, and Jinho Kim. Social network community detection using strongly connected components. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 596–604. Springer, 2014.
- [329] Ruiyun Li, Sen Pei, Bin Chen, Yimeng Song, Tao Zhang, Wan Yang, and Jeffrey Shaman. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (sars-cov-2). *Science*, 368(6490):489–493, 2020.
- [330] Shusong Li and Shiyong Zhang. The opinion dynamics and bounded confidence model on flocking movement world. In *2010 IEEE International Conference on Intelligent Systems and Knowledge Engineering*, pages 355–259, Hangzhou, China, 2011. IEEE.
- [331] Z Li, V Hayse, I Hlohowskyj, K Smith, and R Smith. Agent-based model for simulation of west nile virus transmission. In *Proceedings of the Agent 2005 Conference on Social Dynamics: Interaction, Reflexivity and Emergence; June 26–28; Chicago, USA*, volume 14, 2005.
- [332] Thomas M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, volume 324. Springer-Verlag Berlin Heidelberg, 1 edition, 1999.
- [333] Quah-chi Liu, Marco AJelli, Alberto Aleta, Stefano Merier, Yamir Moreno, and Alessandro Vespignani. Measurability of the epidemic reproduction number in data-driven contact networks. *PNAS*, 115(50):12680–12685, 11 2018.
- [334] Xiangsen Liu and Binxiang Dai. Qualitative and bifurcation analysis of an sir epidemic model with saturated treatment function and nonlinear pulse vaccination. *Discrete Dynamics in Nature and Society*, 2016, 2016.

- [335] Ying Liu, Albert A Gayle, Annelies Wilder-Smith, and Joacim Rocklöv. The reproductive number of covid-19 is higher compared to sars coronavirus. *Journal of travel medicine*, 2020.
- [336] Fabián Aguirre López, Paolo Barucca, Mathilde Fekom, and Anthony CC Coolen. Exactly solvable random graph ensemble with extensively many short cycles. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085101, 2018.
- [337] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025, 2011.
- [338] Linyuan Lü, Duan-Bing Chen, and Tao Zhou. The small world yields the most effective information spreading. *New Journal of Physics*, 13(12):123005, 2011.
- [339] Xin Lu and Christa Brelsford. Network structure and community evolution on twitter: human behavior change in response to the 2011 japanese earthquake and tsunami. *Scientific reports*, 4:6773, 2014.
- [340] Jonas F Ludvigsson. Systematic review of covid-19 in children shows milder cases and a better prognosis than adults. *Acta Paediatrica*, 109(6):1088–1095, 2020.
- [341] Dih-Ling Luh, Cheng-Chieh Liu, Yun-Ru Luo, and Szu-Chieh Chen. Economic cost and burden of dengue during epidemics and non-epidemic years in taiwan. *Journal of Infection and Public Health*, 11(2):215 – 223, 2018.
- [342] Shuangling Luo, Haoxiang Xia, and Borui Yin. Continuous opinion dynamics on an adaptive coupled random network. *Advances in Complex Systems*, 17(03n04):1450012, 2014.
- [343] AA Lushnikov. Exactly solvable model of a coalescing random graph. *Physical Review E*, 91(2):022119, 2015.
- [344] Noni E. MacDonald and the SAGE Working Group on Vaccine Hesitancy. Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, 33(34):4161–4164, 2015.
- [345] Elisabeth Mahase. COVID-19: What is the evidence for cloth masks?, 2020. Accessed: 18 May 2020.
- [346] Kalyeena Makortoff. Astrazeneca doubles capacity for potential COVID-19 vaccine to 2bn doses, Jun 2020. Accessed: 7 June 2020.
- [347] Iman Mallakpour and Gabriele Villarini. A simulation study to examine the sensitivity of the pettitt test to detect abrupt changes in mean. *Hydrological Sciences Journal*, 61(2):245–254, 2016.

- [348] Ryan Martin, Jianchun Zhang, and Chuanhai Liu. Dempster-Shafer theory and statistical inference with weak beliefs. *Statistical Science*, 25(1):72–87, 2010.
- [349] André CR Martins. Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics C*, 19(04):617–624, 2008.
- [350] Joel N. Maslow. The cost and challenge of vaccine development for emerging and emergent infectious diseases. *The Lancet Global Health*, 6(12):e1266 – e1267, 12 2018.
- [351] Naoki Masuda, Michiko Sakaki, Takahiro Ezaki, and Takamitsu Watanabe. Clustering coefficients for correlation networks. *Frontiers in Neuroinformatics*, 12, Mar 2018.
- [352] Joan T Matamalas, Alex Arenas, and Sergio Gómez. Effective approach to epidemic containment using link equations in complex networks. *Science advances*, 4(12):eaau4212, 2018.
- [353] Dileep Mavalankar, Dipti Govil, Neha Trivedi, and Vinubhai Patel. Prevalence of various symptoms and cost of treatment during the chikungunya epidemic in ahmedabad, gujarat, india, in 2006. *Dengue Bulletin*, pages 46–53, 12 2010.
- [354] I McLachlan, G Marion, IJ McKendrick, T Porphyre, IG Handel, and BM deC Bronsvort. Endemic foot and mouth disease: pastoral in-herd disease dynamics in sub-saharan africa. *Scientific reports*, 9(1):1–12, 2019.
- [355] John M McLaughlin, Justin J. McGinnis, Litjen Tan, Annette Mercatante, and Joseph Fortuna. Infection dynamics on scale-free networks. *The Journal of Primary Prevention*, 36(4):259–273, 2015.
- [356] Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.
- [357] Itamar Megiddo, Abigail R Colson, Arindam Nandi, Susmita Chatterjee, Shankar Prinja, Ajay Khera, and Ramanan Laxminarayan. Analysis of the universal immunization programme and introduction of a rotavirus vaccine in india with indiasim. *Vaccine*, 32:A151–A161, 2014.
- [358] AK Melikov, ZT Ai, and DG Markov. Intermittent occupancy combined with ventilation: An efficient strategy for the reduction of airborne transmission indoors. *Science of The Total Environment*, page 140908, 2020.
- [359] Hanna Melnick, Linda Darling-Hammond, M Leung, C Yun, A Schachner, S Plasencia, and N Ondrasek. Reopening schools in the context of covid-19: Health and safety guidelines from other countries. *Learning Policy Institute*, 2020.

- [360] Sandro Meloni, Nicola Perra, Alex Arenas, Sergio Gómez, Yamir Moreno, and Alessandro Vespignani. Modeling human mobility responses to the large-scale spreading of infectious diseases. *Scientific Reports*, 1(62), 8 2011.
- [361] Filippo Menczer. Fake online news spreads through social echo chambers. *Scientific American*, 28, 2016.
- [362] J. R. Meszaros, D. A. Asch, J. Baron, J. C. Herschey, H. Kunreuther, and J. Schwartz-Buzaglo. Cognitive processes and the decisions of some parents to forego pertussis vaccination for their children. *Journal of Clinical Epidemiology*, 49(6):697–703, 1996.
- [363] Lauren Meyers. Contact network epidemiology: Bond percolation applied to infectious disease prediction and control. *Bulletin of the American Mathematical Society*, 44(1):63–86, 2007.
- [364] P. van Mieghem. Epidemic phase transition of the sis type in networks. *Europhysics Letters*, 97(4):48004:1–5, 2 2012.
- [365] Joel C Miller. Percolation and epidemics in random clustered networks. *Physical Review E*, 80(2):020901, 2009.
- [366] Patti J. Miller, Sandra Landry, Mary Ann Searcy, Ella Hunt, and Richard P. Wenzel. Cost of varicella epidemic. *Pediatrics*, 75(5):989 – 990, 5 1985.
- [367] Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42, 2007.
- [368] Lewis Mitchell, Morgan R Frank, Kameron Decker Harris, Peter Sheridan Dodds, and Christopher M Danforth. The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PloS one*, 8(5), 2013.
- [369] Cristopher Moore and M. E. J. Newman. Epidemics and percolation in small-world networks. *Physical Review E*, 61(5):5678, 5 2000.
- [370] Thomas W Moore, Patrick D Finley, Benjamin J Apelberg, Bridget K Ambrose, Nancy S Brodsky, Theresa J Brown, Corinne Husten, and Robert J Glass. An opinion-driven behavioral dynamics model for addictive behaviors. *The European Physical Journal B*, 88(4):95, 2015.
- [371] P.A.P. Moran. The interpretation of statistical maps. *Journal of the Royal Statistical Society Series B (Methodological)*, 10(2):243–251, 1948.
- [372] Nicole Mortillaro. Science around coronavirus is still evolving, and that has left some feeling frustrated and skeptical — cbc news, May 2020. Accessed: 18 May 2020.

- [373] Mosaic. How france is persuading its citizens to get vaccinated, June 2019. [Online; posted 14-October-2019].
- [374] Farzad Mostashari, Martin Kulldorff, Jessica J Hartman, James R Miller, and Varuni Kulasekera. Dead bird clusters as an early warning system for west nile virus activity. *Emerging infectious diseases*, 9(6):641, 2003.
- [375] Matt Motta, Dominik Stecula, and Christina Farhart. How right-leaning media coverage of COVID-19 facilitated the spread of misinformation in the early stages of the pandemic in the us. *Canadian Journal of Political Science/Revue canadienne de science politique*, pages 1–9, 2020.
- [376] Mehdi Moussaïd, Niriaska Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PloS one*, 5(4):e10047, 2010.
- [377] Sudip Mukherjee and Arnab Chatterjee. Disorder-induced phase transition in an opinion dynamics model: Results in two and three dimensions. *Physical Review E*, 94(062317):1–5, 12 2016.
- [378] Simon H. Murch, Andrew Anthony, David H. Casson, Mohsin Malik, Mark Berelowitz, Amar P. Dhillon, Michael A. Thomson, Alan Valentine, Susan E. Davies, and John A. Walker-Smith. Retraction of an interpretation. *Lancet*, 363(9411):750, 2004.
- [379] Jillian Murray and Adam L. Cohen. *Infectious Disease Surveillance*, pages 222–229. Elsevier, 2 edition, 2017.
- [380] David G Myers and George D Bishop. Discussion effects on racial attitudes. *Science*, 169(3947):778–779, 1970.
- [381] M.F. Myers, D.J. Rogers, J. Cox, A. Flahault, and S.I. Hay. Forecasting disease risk for increased epidemic preparedness in public health. *Advanced Parasitology*, 47:309–330, 2000.
- [382] Seth A Myers, Aneesh Sharma, Pankaj Gupta, and Jimmy Lin. Information network or social network? the structure of the twitter follow graph. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 493–498, 2014.
- [383] Xiaoli Nan and Kelly Madden. Hpv vaccine information in the blogosphere: how positive and negative blogs influence vaccine-related risk perceptions, attitudes, and behavioral intentions. *Health communication*, 27(8):829–836, 2012.
- [384] Ingemar Näsell. A new look at the critical community size for childhood infections. *Theoretical Population Biology*, 67(3):203–216, 5 67.

- [385] National Research Council. *Condensed-Matter Physics*, chapter 3, pages 75–94. The National Academies Press, Washington, D.C., 1986.
- [386] M. E. Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(016128), 2002.
- [387] M. E. J. Newman and D. J. Watts. Scaling and percolation in the small-world network model. *Physics Revelation E*, 60(6):7332–7342, 12 1999.
- [388] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1):2566–2572, 2002.
- [389] Mark EJ Newman. Clustering and preferential attachment in growing networks. *Physical review E*, 64(2):025102, 2001.
- [390] Mark EJ Newman. The structure of scientific collaboration networks. *Proceedings of the national academy of sciences*, 98(2):404–409, 2001.
- [391] Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.
- [392] Mark EJ Newman et al. Random graphs as models of networks. *Handbook of graphs and networks*, 1:35–68, 2003.
- [393] Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.
- [394] M.E.J. Newman. Mixing patterns in networks. *Physical Review E*, page 026126, 2003.
- [395] N.E.J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20):1–4, 2002.
- [396] CBC News. Toronto public health urges city’s largest school board to keep class sizes down. <https://www.cbc.ca/news/canada/toronto/toronto-public-health-school-guidance-1.5677739>, 2020; accessed August 7, 2020.
- [397] Huyen Nguyen, Charles Huyi, and Paul Warren. Mitigating the spread of fake news by identifying and disrupting echo chambers. https://www.k4all.org/wp-content/uploads/2017/09/WPOC2017_paper_10.pdf, 2017. Accessed 26 April 2020.
- [398] C Thi Nguyen Associate Professor of Philosophy. The problem of living inside echo chambers, Sep 2019.
- [399] A Nicoll, D Elliman, and NT Begg. Immunisation: causes of failure and strategies and tactics for success. *BMJ: British Medical Journal*, 299(6703):808, 1989.

- [400] Hiroshi Nishiura, Natalie M Linton, and Andrei R Akhmetzhanov. Serial interval of novel coronavirus (2019-ncov) infections. *medRxiv*, 2020.
- [401] Andrzej Nowak, Jacek Szamrej, and Bibb Latané. From private attitude to public opinion. *Psychological Review*, 97(3):362–376, 1990.
- [402] Brendan Nyhan and Jason Reifler. When corrections fail: The persistence of political misperceptions. *Political Behaviour*, 32(2):303–330, 2010.
- [403] Godfrey P Oakley and Richard B Johnston. Balancing benefits and harms in public health prevention programmes mandated by governments. *British Medical Journal*, 329(7456):41–43, 2004.
- [404] Eamon B. O’Dea, Andrew W. Park, and John M. Drake. Estimating the distance to an epidemic threshold. *Journal of the Royal Society Interface*, 15(20180034):1–10, 2018.
- [405] Institute of Medicine (US) Vaccine Safety Forum. Risk communication and vaccination: Summary of a workshop. Risk Perception and Decision making, 1997.
- [406] Government of Ontario. Covid-19: reopening child care centres. <http://www.edu.gov.on.ca/childcare/child-care-re-opening-operational-guidance.pdf>, 2020; accessed August 7, 2020.
- [407] U.S. Office of Personnel Management. Fact Sheet: The Use of Flexible Work Schedules in Response to Coronavirus Disease 2019 (COVID-19), 2020.
- [408] Saad B Omer, Daniel A Salmon, Walter A Orenstein, M Patricia Dehart, and Neal Halsey. Vaccine refusal, mandatory immunization, and the risks of vaccine-preventable diseases. *New England Journal of Medicine*, 360(19):1981–1988, 2009.
- [409] National Research Council (US) Committee on Achieving Sustainable Global Capacity for Surveillance and Response to Emerging Diseases of Zoonotic Origin. *Diseases in Humans: Early Warning Systems*, chapter 4, pages 53–68. National Academies Press (US), 2008.
- [410] The SAGE Working Group on Vaccine Hesitancy. Report of the SAGE working group on vaccine hesitancy. Technical report, World Health Organisation, October 2014.
- [411] Graziano Onder, Giovanni Rezza, and Silvio Brusaferro. Case-fatality rate and characteristics of patients dying in relation to COVID-19 in italy. *Jama*, 323(18):1775–1776, 2020.
- [412] Tore Opsahl and Pietro Panzarasa. Clustering in weighted networks. *Social networks*, 31(2):155–163, 2009.

- [413] Tamer Oraby, Vivek Thampi, and Chris T. Bauch. The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proceedings of the Royal Society B*, 281(20133172), 2014.
- [414]Carolynn Orbann, Lisa Sattenspiel, Erin Miller, and Jessica Dimka. Defining epidemics in computer simulation models: How do definitions influence conclusions? *Epidemics*, 19:24–32, 2017.
- [415] Suzanne M. O’Regan and John M. Drake. Theory of early warning signals of disease emergence and leading indicators of elimination. *Theoretical Ecology*, 6(3):333–357, 8 2013.
- [416] Suzanne M. O’Regan, Jonathan W. Lillie, and John M. Drake. Leading indicators of mosquito-borne disease elimination. *Theoretical Ecology*, 9:269–286, 2016.
- [417] World Health Organisation. Coronavirus Disease (COVID-19) Dashboard.
- [418] World Health Organization. WHO announces COVID-19 outbreak a pandemic. *World Health Organization*, Mar 2020. Assessed: 17 May 2020.
- [419] Günce Keziban Orman and Vincent Labatut. A comparison of community detection algorithms on artificial networks. In *International Conference on Discovery Science*, pages 242–256. Springer, 2009.
- [420] S. Orozco-Fuentes, G. Griffiths, M.J. Holmes, R. Ettelaie, J. Smith, A.W. Baggaley, and N.G. Parker. Early warning signals in plant disease outbreaks. *Ecological Modelling*, 393:12–19, 2 2019.
- [421] Katie Orr. Anti-vaccination activists join stay-at-home order protesters, May 2020. Accessed: 18 May 2020.
- [422] Michael T. Osterholm and Craig W. Hedberg. Epidemiologic principles. In John E. Bennett, Raphael Dolin, and Martin J. Blaser, editors, *Mandell, Douglas, and Bennett’s Principles and Practice of Infectious Diseases*, volume 1, chapter 13, pages 146–157. Elsevier Inc., 8 edition, 2014.
- [423] Sachiko Ozawa and Meghan L Stack. Public trust and vaccine acceptance-international perspectives. *Human vaccines & immunotherapeutics*, 9(8):1774–1778, 2013.
- [424] Pandemic Influenza Outbreak Research Modelling Team (Pan-InfORM et al. Modelling an influenza pandemic: A guide for the perplexed. *Cmaj*, 181(3-4):171–173, 2009.
- [425] A Demetri Pananos, Thomas M Bury, Clara Wang, Justin Schonfeld, Sharada P Mohanty, Brendan Nyhan, Marcel Salathé, and Chris T Bauch. Critical dynamics in population vaccinating behavior. *Proceedings of the National Academy of Sciences*, 114(52):13762–13767, 2017.

- [426] Krešo Pandžić, Mira Kobold, Dijana Oskoruš, Božidar Biondić, Ranko Biondić, Ognjen Bonacci, Tanja Likso, and Oliver Curić. Standard normal homogeneity test as a tool to detect change points in climate-related river discharge variation: case study of the kupa river basin. *Hydrological Sciences Journal*, 65(2):227–241, 2020.
- [427] Eli Pariser. *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.
- [428] Sang Woo Park, Benjamin M Bolker, David Champredon, David JD Earn, Michael Li, Joshua S Weitz, Bryan T Grenfell, and Jonathan Dushoff. Reconciling early-outbreak estimates of the basic reproductive number and its uncertainty: framework and applications to the novel coronavirus (sars-cov-2) outbreak. *MedRxiv*, 2020.
- [429] Cameron Parker. The influence relation for ternary voting games. *Games and Economic Behavior*, 75(2):867–881, 2012.
- [430] H Van Dyke Parunak, Robert Savit, and Rick L Riolo. Agent-based modeling vs. equation-based modeling: A case study and users’ guide. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pages 10–25. Springer, 1998.
- [431] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of modern physics*, 87(3):925, 2015.
- [432] R. K. Pathria. *Statistical mechanics*. Butterworth-Heinemann, Oxford ;, 2nd ed. edition, 1996.
- [433] Padmavathi Patlolla, Vandana Gunupudi, Armin R Mikler, and Roy T Jacob. Agent-based simulation tools in computational epidemiology. In *International Workshop on Innovative Internet Community Systems*, pages 212–223. Springer, 2004.
- [434] Stacy Patterson and Bassam Bamieh. Interaction-driven opinion dynamics on online social networks. In *Proceesings of the First Workshop on Social Media Analytics*, pages 98–105, 2010.
- [435] Michael Paty. D’alembert et les probabilités. In Roshdi Rashed, editor, *Sciences à l’époque del la Révolution française*, pages 205–265. Recherches historiques, 1988. HAL ID: alshs-00004289.
- [436] Michael J. Paul, Mark Dredze, and David Broniatowski. Twitter improves influenza forecasting. *PLOS Currents Outbreaks*, 10 2014.
- [437] L. Pellis, N.M. Fergusson, and C. Fraser. Threshold parameters for a model of epidemic spread among households and workplaces. *Journal of the Royal Society Interface*, 6(40):979–987, 2 2009.

- [438] Marzio Pennisi, Giulia Russo, Silvia Ravalli, and Francesco Pappalardo. Combining agent based-models and virtual screening techniques to predict the best citrus-derived vaccine adjuvants against human papilloma virus. *BMC bioinformatics*, 18(16):544, 2017.
- [439] Alex "Sandy" Pentland. Beyond the echo chamber. *Harvard Business Review*, Nov 2013.
- [440] Liliana Perez and Suzana Dragicevic. An agent-based approach for modeling dynamics of contagious disease spread. *International journal of health geographics*, 8(1):50, 2009.
- [441] A.N. Pettitt. A non-parametric approach to the change-point problem. *Applied Statistics*, 28:126–135, 1979.
- [442] Arun G Phadke and James S Thorp. *Computer relaying for power systems*. John Wiley & Sons, 2009.
- [443] Varun K. Phadke, Robert A. Bednarczyk, Daniel A. Salmon, and Saad 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.
- [444] Tuan Q Phan and Edoardo M Airoidi. A natural experiment of social network formation and dynamics. *Proceedings of the National Academy of Sciences*, 112(21):6595–6600, 2015.
- [445] Brendon Phillips, Madhur Anand, and Chris T Bauch. Spatial early warning signals of social and epidemiological tipping points in a coupled behaviour-disease network. *Scientific reports*, 10(1):1–12, 2020.
- [446] Brendon Phillips and Chris T. Bauch. Echo chambers as early warning signals of widespread vaccine refusal in social-epidemiological networks. 2020.
- [447] Brendon Phillips, Dillon Browne, Madhur Anand, and Chris Bauch. Model-based projections for COVID-19 outbreak size and student-days lost to closure in ontario childcare centres and primary schools. 2020.
- [448] André Picard. Advice on masks is changing as coronavirus knowledge evolves, Apr 2020. Accessed: 14 May 2020.
- [449] Joseph M Pierre. Anti-vaccine and COVID-19 conspiracy theories: A perfect storm, May 2020. Accessed: 7 June 2020.
- [450] CA Piña-García, Carlos Gershenson, and J Mario Siqueiros-García. Towards a standard sampling methodology on online social networks: Collecting global trends on twitter. *Applied network science*, 1(1):3, 2016.

- [451] CA Piña-García and Dongbing Gu. Spiraling facebook: an alternative metropolis–hastings random walk using a spiral proposal distribution. *Social Network Analysis and Mining*, 3(4):1403–1415, 2013.
- [452] Marcelo A. Pires, André L. Oestereich, and Nuno Crokidakis. Sudden transitions in coupled opinion and epidemic dynamics with vaccination. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(053407), 2018.
- [453] Richard E. Plant. *The R Programming Environment*. CRC Press, 2018.
- [454] Thorsten Pohlert. *trend: Non-Parametric Trend Tests and Change-Point Detection*, 2018. R package version 1.1.1.
- [455] Piero Poletti, Bruno Caprile, Marco Ajelli, Andrea Pugliese, and Stefano Merler. Spontaneous behavioural changes in response to epidemics. *Journal of Theoretical Biology*, 260(1), 9 2009.
- [456] Philip M Polgreen, Yiling Chen, David M Pennock, Forrest D Nelson, and Robert A Weinstein. Using internet searches for influenza surveillance. *Clinical infectious diseases*, 47(11):1443–1448, 2008.
- [457] Yves Pomeau and Martine Le Berre. Critical speed-up vs critical slow-down: a new kind of relaxation oscillation with application to stick-slip phenomena. *arXiv preprint arXiv:1107.3331*, 2011.
- [458] Kiesha Prem, Alex R Cook, and Mark Jit. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLoS computational biology*, 13(9):e1005697, 2017.
- [459] Heather Prime, Mark Wade, and Dillon T Browne. Risk and resilience in family well-being during the covid-19 pandemic. *American Psychologist*, 2020.
- [460] Public Health Ontario. Ontario covid-19 data tool. <https://www.publichealthontario.ca/en/data-and-analysis/infectious-disease/covid-19-data-surveillance/covid-19-data-tool>, accessed June 10, 2020.
- [461] Walter Quattrociocchi. Part 2-social and political challenges: 2.1 western democracy in crisis. In *World Economic Forum*, 2017.
- [462] Walter Quattrociocchi, Guido Caldarelli, and Antonio Scala. Opinion dynamics on interacting networks: media competition and social influence. *Scientific reports*, 4:4938, 2014.
- [463] Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. Echo chambers on facebook. *Available at SSRN 2795110*, 2016.

- [464] Amira Rachah et al. Analysis, simulation and optimal control of a seir model for ebola virus with demographic effects. *Communications Faculty of Sciences University of Ankara Series A1 Mathematics and Statistics*, 67(1):179–197, 2018.
- [465] Filippo Radicchi, Claudio Castellano, Federico Cecconi, Vittorio Loreto, and Domenico Parisi. Defining and identifying communities in networks. *Proceedings of the national academy of sciences*, 101(9):2658–2663, 2004.
- [466] Hazhir Rahmandad and John Sterman. Homogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management Science*, 54(5):998–1014, 2008.
- [467] Amol Rajan. Do digital echo chambers exist? *BBC News*, Mar 2019.
- [468] Gabriele Ranco, Ilaria Bordino, Giacomo Borgetti, Guido Caldarelli, Fabrizio Lillo, and Michele Treccani. Coupling news sentiment with web browsing data improves prediction of intra-day price dynamics. *PLOS ONE*, 11(1):1–14, 01 2016.
- [469] T. S. Sathyanarayana Rao and Chittaranjan Andrade. The mmr vaccine and autism: Sensation, refutation, retraction, and fraud. *Indian Journal of Psychiatry*, 53(2):95–96, 2011.
- [470] Phil Reed. Are echo chambers a threat to intellectual freedom? *Psychology Today*, Mar 2019.
- [471] Hui Ren and David Watts. Early warning signals for critical transitions in power systems. *Electric Power Systems Research*, 124:173–180, 2015.
- [472] Craig W Reynolds. *Flocks, herds and schools: A distributed behavioral model*, volume 21. ACM, 1987.
- [473] David Robson. The myth of the online echo chamber. *BBC Future*, Apr 2018.
- [474] Ronald Rogers. A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology Interdisciplinary and Applied*, 91(1):93–114, 1975.
- [475] Anne Rooney. *Cheer Up, It May Never Happen*, chapter 7, pages 170–171. Rosen Publishing Group, 1 edition, 2013.
- [476] Suman Deb Roy and Wenjun Zeng. *Social multimedia signals*. Springer, 2014.
- [477] Faryad Darabi Sahneh, Fahmida N. Chowdhury, and Caterina M. Scoglio. On the existence of a threshold for preventative behavioral responses to suppress epidemic spreading. *Scientific Reports*, 2(632), 9 2012.

- [478] Marcel Salathé and Sebastian Bonhoeffer. The effect of opinion clustering on disease outbreaks. *Journal of The Royal Society Interface*, 5:1505–1508, 2008.
- [479] Marcel Salathé and James H Jones. Dynamics and control of diseases in networks with community structure. *PLoS computational biology*, 6(4):e1000736, 2010.
- [480] Marcel Salathé and Shashank Khandelwal. Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control. *PLoS computational biology*, 7(10):e1002199, 2011.
- [481] LM SANDAR. Percolation on heterogeneous networks as a model for epidemics. *Mathematical Biosciences*, 180:293, 2002.
- [482] Joshua L Santarpia, Vicki L Herrera, Danielle N Rivera, Shanna Ratnesar-Shumate, Paul W Denton, Jacob WS Martens, Ying Fang, Nicholas Conoan, Michael V Callahan, James V Lawler, et al. The infectious nature of patient-generated sars-cov-2 aerosol. *medRxiv*, 2020.
- [483] Nadir Sari and Emmanuelle Augeraud-Véron. Periodic orbits of a seasonal sis epidemic model with migration. *Journal of Mathematical Analysis and Applications*, 423(2):1849–1866, 2015.
- [484] M. Scheffer, J. Bascompte, W. Brock, V. Brovkin, S. Carpenter, V. Dakos, H. Held, E. van Nes, M. Rietkerk, and G. Sugihara. Early-warning signals for critical transitions. *Nature*, 461(3), 2009.
- [485] Marten Scheffer. *Critical Transitions in Nature and Society*. Princeton University Press, Princeton, New Jersey, 2009.
- [486] Martin Scheffer, Stephen Carpenter, Timothy M. Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan De Koppel, Ingrid A. Van De Leemput, Simon A. Levin, Egbert H. Van Nes, Mercedes Pascual, and John Vandermeer. Anticipating critical transitions. *Science*, 338(6105):344–348, 2012.
- [487] Thomas C. Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1:143–186, 1971.
- [488] David Schkade, Cass R Sunstein, and Daniel Kahneman. Deliberating about dollars: The severity shift. *Colum. L. Rev.*, 100:1139, 2000.
- [489] Ana Lucía Schmidt, Fabiana Zollo, Michela Del Vicario, Alessandro Bessi, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. Anatomy of news consumption on facebook. *Proceedings of the National Academy of Sciences*, 114(12):3035–3039, 2017.

- [490] Ana Lucía Schmidt, Fabiana Zollo, Antonio Scala, Cornelia Betsch, and Walter Quattrocchi. Polarization of the vaccination debate on facebook. *Vaccine*, 36(25):3606–3612, 2018.
- [491] Christopher Seneca. How to break out of your social media echo chamber. *Wired*, Sep 2020.
- [492] M Ángeles Serrano and Marián Boguná. Percolation and epidemic thresholds in clustered networks. *Physical review letters*, 97(8):088701, 2006.
- [493] Paul M. Severns, Laura K. Estep, Kathryn E. Sackett, and Christophe C. Mundt. Degree of host susceptibility on the initial disease outbreak influence subsequent epidemic spread. *Journal of Applied Ecology*, 51:1622–1630, 2014.
- [494] Jian Shang, Gang Ye, Ke Shi, Yushun Wan, Chuming Luo, Hideki Aihara, Qibin Geng, Ashley Auerbach, and Fang Li. Structural basis of receptor recognition by sars-cov-2. *Nature*, pages 1–4, 2020.
- [495] Siyao Shao, Dezhi Zhou, Ruichen He, Jiaqi Li, Shufan Zou, Kevin Mallery, Santosh Kumar, Suo Yang, and Jiarong Hong. Assessment of airborne transmission potential of covid-19 by asymptomatic individuals under different practical settings. *arXiv preprint arXiv:2007.03645*, 2020.
- [496] Shixin (Cindy) Shen and Vinita Dubey. Addressing vaccine hesitancy. *Canadian Family Physician*, 65(3):175–181, 2019.
- [497] Feng Shi, Misha Teplitskiy, Eamon Duede, and James A Evans. The wisdom of polarized crowds. *Nature human behaviour*, 3(4):329, 2019.
- [498] SickKids. Covid-19: Guidance for school reopening. <http://www.sickkids.ca/pdfs/about-sickkids/81407-covid19-recommendations-for-school-reopening-sickkids.pdf>, 2020; accessed August 7, 2020.
- [499] Constantinos Siettos, Cleo Anastassopoulou, Lucia Russo, Christos Grigoras, and Eleftherios Mylonakis. Modeling the 2014 ebola virus epidemic—agent-based simulations, temporal analysis and future predictions for liberia and sierra leone. *PLoS currents*, 7, 2015.
- [500] Alessio Signorini, Alberto Maria Segre, and Philip M Polgreen. The use of twitter to track levels of disease activity and public concern in the us during the influenza a h1n1 pandemic. *PloS one*, 6(5):e19467, 2011.
- [501] Martin Silberberg. *Principles of general chemistry*. McGraw-Hill Education, 2012.
- [502] Alina Sirbu, Vittorio Loreto, Vito D. P. Servedio, and Francesca Tria. Opinion dynamics: Models, extensions and external effects. *Participatory Sensing, Opinions and Collective Awareness*, pages 363–401, 5 2016.

- [503] Jackson Andrew Smith, Jean de Dieu Basabose, Dillon Thomas Browne, C Psych, and Michael Stephenson. Family medicine with refugee newcomers during the covid-19 crisis. *Journal of the American Board of Family Medicine*, 2020.
- [504] Louise E. Smith, John Weinman, Richard Amlôt, Jenny Yiend, and G. James Rubin. Parental expectation of side effects following vaccination is self-fulfilling: A prospective cohort study. *Annals of Behavioral Medicine*, 53(3):267 – 282, 6 2019.
- [505] P. J. Smith, S. G Humiston, E. K. Marcuse, Z. Zhao, C. G. Dorrell, C. Howes, and B. Hibbs. Parental delay or refusal of vaccine doses, childhood vaccination coverage at 24 months of age, and the Health Belief Model. *Public Health Reports*, 126:135–146, 2011.
- [506] Tara C. Smith. Vaccine rejection and hesitancy: A review and call to action. *Open Forum Infectious Diseases*, 4(3), 7 2017.
- [507] Tara C. Smith. Coronavirus vaccine and quarantine protesters in america form an unholy COVID-19 alliance, May 2020. Accessed: 18 May 2020.
- [508] Elisa J. Sobo. Social cultivation of vaccine refusal and delay among Waldorf (Steiner) school parents. *Medical Anthropology Quarterly*, 29(3):381 – 399, 4 2015.
- [509] Man-Koumba Soumahoro, Pierre-Yves Boelle, Bernard-Alex Gaüzere, Kokuvi Atsou, Camille Pelat, Bruno Lambert, Guy La Ruche, Marc Gastellu-Etchegorry, Philippe Renault, Marianne Sarazin, Yazdan Yazdanpanah, Antoine Flahault, Denis Malvy, and Thomas Hanslik. The chikungunya epidemic on la réunion island in 2005–2006: A cost-of-illness study. *PLOS Neglected Tropical Diseases*, 5(6):1–9, 06 2011.
- [510] A. Spreco and T Timpka. Algorithms for detecting and predicting influenza outbreaks: metanarrative review of prospective evaluation. *BMJ Open*, 6:e010683, 2016.
- [511] V. Sriluckshmy and Ipsita Mandal. Critical scaling of the mutual information in two-dimensional disordered ising models. *Journal of Statistical Mechanics: Theory and Experiment*, 2018, 4 2018.
- [512] NPR Staff. The reason your feed became an echo chamber - and what to do about it. *NPR*, Jul 2016.
- [513] The World staff. World faces risk of 'vaccine nationalism' in COVID-19 fight, says cepi chair, May 2020. Accessed: 7 June 2020.
- [514] The Star. Thousands sign petition asking ontario to reduce class sizes for elementary school students this fall., 2020; accessed August 7, 2020.

- [515] Chen Stein-Zamir, Nitza Abramson, Hanna Shoob, Erez Libal, Menachem Bitan, Tanya Cardash, Refael Cayam, and Ian Miskin. A large covid-19 outbreak in a high school 10 days after schools' reopening, israel, may 2020. *Eurosurveillance*, 25(29):2001352, 2020.
- [516] J Stover, G Garnett, S Seitz, and S Forsythe. *The epidemiological impact of an HIV/AIDS vaccine in developing countries. World Bank Policy Research Discussion Paper, no. 2811*. Development Research Department, World Bank Washington, DC, 2002.
- [517] Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Application To Physics, Biology, Chemistry, And Engineering*. Westview Press, 2 edition, 7 2014.
- [518] Cass R Sunstein. The law of group polarization. *Journal of political philosophy*, 10(2):175–195, 2002.
- [519] Cass R Sunstein. *Going to extremes: How like minds unite and divide*. Oxford University Press, 2009.
- [520] Supreme Court of Canada. An Act respecting certain measures in response to COVID-19. *COVID-19 Emergency Response Act*, S.C. 2020(5):1–52, 2020.
- [521] James Surowiecki. The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business. *Economies, Societies and Nations*, 296, 2004.
- [522] Liz Szabo. How anti-vaccine activists are using COVID-19 to boost their movement: Spectrum: Autism research news, Apr 2020. Accessed: 18 May 2020.
- [523] Alina Sîrbu, Vittorio Loreto, Vito D.P. Servedio, and Francesca Tria. *Opinion Dynamics: Models, Extensions and External Effects*, pages 363–401. Springer, Cham, Switzerland, 05 2017.
- [524] Alireza Tahbaz-Salehi and Ali Jadbabaie. A necessary and sufficient condition for consensus over random networks. *IEEE Transactions on Automatic Control*, 53(3):791–795, 2008.
- [525] Zhou Tao, Fu Zhongqian, and Wang Binghong. Epidemic dynamics on complex networks. *Progress in Natural Science*, 16(5):452–457, 2006.
- [526] C.J. Tessone, R. Toral, P. Amengual, Wio. H.S., and M. San Miguel. Neighbourhood models of minority opinion spreading. *Europeah Physics Journal*, 2004.
- [527] Nicholas A Thapen and Moustafa M Ghanem. Towards passive political opinion polling using twitter. In *SMA@ BCS-SGAI*, pages 19–34, 2013.
- [528] The Editors of The Lancet. Retraction—Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children. *Lancet*, 375(9713):445, 2010.

- [529] Lauren Tindale, Michelle Coombe, Jessica E Stockdale, Emma Garlock, Wing Yin Venus Lau, Manu Saraswat, Yen-Hsiang Brian Lee, Louxin Zhang, Dongxuan Chen, Jacco Wallinga, et al. Transmission interval estimates suggest pre-symptomatic spread of COVID-19. *medRxiv*, 2020.
- [530] I-Hsien Ting and Hui-Ju Wu. Web mining techniques for on-line social networks analysis: An overview. In *Web Mining Applications in E-commerce and E-services*, pages 169–179. Springer, 2009.
- [531] Mathew Titus and James Watson. Critical speeding up as an early warning signal of stochastic regime shifts. *Theoretical Ecology*, pages 1–9, 2020.
- [532] Isaac Todhunter. *D’Alembert*, chapter 8, pages 265–275. Macmillan and Co., 1865. Original printed in Cambridge, owned by the Bayerische Staatsbibliothek München, and digitised by Google Inc.
- [533] Pritish K. Tosh, Thomas G. Boyce, and Gregory A. Poland. Flu myths: Dispelling the myths associated with live attenuated influenza vaccine. *Mayo Clinic Proceedings*, 83(3):77–84, 1 2008.
- [534] Melissa Tracy, Magdalena Cerdá, and Katherine M Keyes. Agent-based modeling in public health: current applications and future directions. *Annual review of public health*, 39:77–94, 2018.
- [535] Barry Trevelyan and Matthew Smallman-Raynor. The spatial structure of epidemic emergence: geographical aspects of poliomyelitis in north-eastern usa, july-october 1916. *Journal of the Royal Statistical Society Series A (Statistics in Society)*, 168(1):701–722, 11 2005.
- [536] Wellcome Trust. Wellcome global monitor. Technical report, Gallup, London NW1 2BE, United Kingdom, 2018.
- [537] Alan Tsang and Kate Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 2016 international conference on autonomous agents & multiagent systems*, pages 368–375, 2016.
- [538] Ryan Tumilty. COVID-19 crisis prompts federal government to look at boosting vaccine production capacity, May 2020. Accessed: 7 June 2020.
- [539] Johan Ugander, Brian Karrer, Lars Backstrom, and Cameron Marlow. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*, 2011.
- [540] Lucas Daniel Valdez, Pablo Alejandro Macri, and Lidia Adriana Braunstein. Temporal percolation of the susceptible network in an epidemic spreading. *PLoS One*, 7(9):e44188, 2012.

- [541] Linda Valeri, Oscar Patterson-Lomba, Yared Gurm, Akweley Ablorh, Jennifer Bobb, F. William Townes, and Guy Harling. Predicting subnational ebola virus disease epidemic dynamics from sociodemographic indicators. *PLOS ONE*, 11(10):e0163544, 2012.
- [542] Srinivasan Venkatramanan, Bryan Lewis, Jiangzhuo Chen, Dave Higdon, Anil Vullikanti, and Madhav Marathe. Using data-driven agent-based models for forecasting emerging infectious diseases. *Epidemics*, 22:43–49, 2018.
- [543] Stéphane Verguet, Mira Johri, Shaun K Morris, Cindy L Gauvreau, Prabhat Jha, and Mark Jit. Controlling measles using supplemental immunization activities: a mathematical model to inform optimal policy. *Vaccine*, 33(10):1291–1296, 2015.
- [544] Zak Vescera. COVID-19: Saskatoon lab explores scaling up vaccine production capacity, May 2020. Accessed: 7 June 2020.
- [545] Eric Vincent and Vincent Dupuis. *Spin glasses: experimental signatures and salient outcomes*, volume 275 of *Springer Series in Materials Science*, chapter 2, pages 31–56. Springer International Publishing, Switzerland, 2018.
- [546] Russell M Viner, Simon J Russell, Helen Croker, Jessica Packer, Joseph Ward, Claire Stansfield, Oliver Mytton, Chris Bonell, and Robert Booy. School closure and management practices during coronavirus outbreaks including covid-19: a rapid systematic review. *The Lancet Child & Adolescent Health*, 2020.
- [547] Mark Wade, Heather Prime, and Dillon T Browne. Why we need longitudinal mental health research with children and youth during (and after) the covid-19 pandemic. *Psychiatry Research*, 2020.
- [548] Liping Wang, Lianmei Jin, Weiyi Xiong, Wenxiao Tu, and Chuchu Ye. *Infectious Disease Surveillance in China*, chapter 2, pages 23–33. Elsevier, 2017.
- [549] Rong Wang, John A Dearing, Peter G Langdon, Enlou Zhang, Xiangdong Yang, Vasilis Dakos, and Marten Scheffer. Flickering gives early warning signals of a critical transition to a eutrophic lake state. *Nature*, 492(7429):419–422, 2012.
- [550] Stanley Wasserman. *Social network analysis : methods and applications*. Structural analysis in the social sciences ; 008. Cambridge University Press, Cambridge [England] ;, 1994.
- [551] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.
- [552] Gerhard Weikum, Nikos Ntarmos, Marc Spaniol, Peter Triantafillou, András A Benczúr, Scott Kirkpatrick, Philippe Rigaux, and Mark Williamson. Longitudinal analytics on web archive data: It’s about time! In *CIDR*, pages 199–202, 2011.

- [553] Gerard Weisbuch. Bounded confidence and social networks. *The European Physical Journal B*, 38(2):339–343, 2004.
- [554] Gérard Weisbuch, Guillaume Deffuant, Frédéric Amblard, and Jean-Pierre Nadal. Meet, discuss, and segregate! *Complexity*, 7(3):55–63, 2002.
- [555] Chad R Wells, Eili Y Klein, and Chris T Bauch. Policy resistance undermines superspreader vaccination strategies for influenza. *PLoS Comput Biol*, 9(3):e1002945, 2013.
- [556] Lilian Weng, Alessandro Flammini, Alessandro Vespignani, and Filippo Menczer. Competition among memes in a world with limited attention. *Scientific reports*, 2:335, 2012.
- [557] United Nations Development Group (UNDG) Western and Central Africa. Socio-economic impact of ebola virus disease in west african countries: A call for national and regional containment, recovery and prevention. Technical report, United Nations Development Group (UNDG) - Western and Central Africa, 2 2015.
- [558] Lander Willem, Sean Stijven, Ekaterina Vladislavleva, Jan Broeckhove, Philippe Beutels, and Niel Hens. Active learning to understand infectious disease models and improve policy making. *PLoS computational biology*, 10(4):e1003563, 2014.
- [559] Lander Willem, Frederik Verelst, Joke Bilcke, Niel Hens, and Philippe Beutels. Lessons from a decade of individual-based models for infectious disease transmission: a systematic review (2006-2015). *BMC infectious diseases*, 17(1):612, 2017.
- [560] Andrew DC Williams, Ian M Hall, G James Rubin, Richard Amlôt, and Steve Leach. An individual-based simulation of pneumonic plague transmission following an outbreak and the significance of intervention compliance. *Epidemics*, 3(2):95–102, 2011.
- [561] Reid Wilson. *Epidemic. Ebola and the Global Scramble to Prevent the Next Killer Outbreak*. Brookings Institution Press, 3 2018.
- [562] C. Wissel. A universal law of the characteristic return time near thresholds. *Oecologia*, 65:101–107, 12 1984.
- [563] Tessa Wong. Coronavirus: Why some countries wear face masks and others don’t, May 2020. Accessed: 16 May 2020.
- [564] Joseph T Wu, Kathy Leung, Mary Bushman, Nishant Kishore, Rene Niehus, Pablo M de Salazar, Benjamin J Cowling, Marc Lipsitch, and Gabriel M Leung. Estimating clinical severity of COVID-19 from the transmission dynamics in wuhan, china. *Nature Medicine*, 26(4):506–510, 2020.
- [565] Qingchu Wu, Xinchu Fu, Michael Small, and Xin-Jian Zu. The impact of awareness on epidemic spreading in networks. *Chaos*, 22:013101, 1 2012.

- [566] Chengyi Xia, Zhishuang Wang, Chunyuan Zheng, Quantong Guo, Yongtang Shi, Matthias Dehmer, and Zengqiang Chen. A new coupled disease-awareness spreading model with mass media on multiplex networks. *Information Sciences*, 471:185–200, 1 2019.
- [567] Shang Xia and Jiming Liu. A belief-based model for characterizing the spread of awareness and its impacts on individuals’vaccination decisions. *Journal of The Royal Society Interface*, 11(94), 2014.
- [568] Jierui Xie, Jeffrey Emenheiser, Matthew Kirby, Sameet Sreenivasan, Boleslaw K. Szymanski, and Gyorgy Korniss. Evolution of opinions on social networks in the presence of competing committed groups. *PLOS ONE*, 7(3), 2012.
- [569] Bo Yang, Jiming Liu, Jianfeng Feng, and Dayou Liu. On modularity of social network communities: The spectral characterization. In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology-Volume 01*, pages 127–133. IEEE Computer Society, 2008.
- [570] Hui Yang, Tim Rogers, and Thilo Gross. Network inoculation: Heteroclinics and phase transitions in an epidemic model. *Chaos*, 26:083116, 8 2016.
- [571] Weizhong Yang, Yajia Lan, Qiao Sun, Wangm Jinfeng, and Zhongjie Li. *Introduction*, chapter 1, pages 3–22. Elsevier, 2017.
- [572] Zhao Yang, René Algesheimer, and Claudio J Tessone. A comparative analysis of community detection algorithms on artificial networks. *Scientific reports*, 6:30750, 2016.
- [573] Julia M Yeomans. *Statistical mechanics of phase transitions*. Clarendon Press, 1992.
- [574] Meredith E. Young, Geoffrey R. Norman, and Karin R. Humphreys. Medicine in the popular press: The influence of the media on perceptions of disease. *PLOS ONE*, 3(10):1–7, 10 2008.
- [575] Yi Yu and Gaoxi Xiao. Influence of random opinion change in complex networks. In *2015 IEEE International Conference on Digital Signal Processing (DSP)*, pages 750–754. IEEE, 2015.
- [576] Lotfi A. Zadeh. A simple view of the Dempster-Shafer theory of evidence and its implication for the rule of combination. *Artificial Intelligence Magazine*, 7(2), 1986.
- [577] Reza Zafarni, Mohammad Ali Abbasi, and Huan Liu. *Social Media Mining: An Introduction*. Cambridge University Press, 4 2018.
- [578] Savvas Zannettou, Barry Bradlyn, Emiliano De Cristofaro, Haewoon Kwak, Michael Sirivianos, Gianluca Stringini, and Jeremy Blackburn. What is gab: A bastion of free speech or an alt-right echo chamber. In *Companion Proceedings of the The Web Conference 2018*, pages 1007–1014, 2018.

- [579] Weituo Zhang, Chjan C. Lim, G. Korniss, and Boleslaw K. Szymanski. Opinion dynamics and influencing on random geometric graphs. *Scientific Reports*, 4(5568):1–9, 2014.
- [580] D.S. Zhao, K Zhao, J Yang, and J.R. Sha. A friendship network based on random and triadic-closure in a fixed community. In *Frontiers in Computer Education: Proceedings of the 2nd International Conference on Frontiers in Computer Education (ICFCE 2014), Wuhan, China, December 24-25, 2014*, page 29. CRC Press, 2015.
- [581] Yi Zheng, Edmund Goh, and Jun Wen. The effects of misleading media reports about COVID-19 on chinese tourists’ mental health: a perspective article. *Anatolia*, 31(2):337–340, 2020.
- [582] Fabiana Zollo. Dealing with digital misinformation: a polarised context of narratives and tribes. *EFSA Journal*, 17:e170720, 2019.
- [583] Fabiana Zollo, Alessandro Bessi, Michela Del Vicario, Antonio Scala, Guido Caldarelli, Louis Shekhtman, Shlomo Havlin, and Walter Quattrociocchi. Debunking in a world of tribes. *PloS one*, 12(7), 2017.
- [584] Lucila G Alvarez Zuzek, Camila Buono, and Lidia A Braunstein. Epidemic spreading and immunization strategy in multiplex networks. In *Journal of Physics: Conference Series*, volume 640, page 012007. IOP Publishing, 2015.

APPENDICES

Appendix A

**Spatial aggregation can indicate
regime shifts**

A.1 \mathcal{I} (Moran's I) and \mathcal{C} (Geary's C) in terms of join counts

We support our previous assertion that the Moran's I (\mathcal{I}) and Geary's C (\mathcal{C}) coefficients are linear combinations of join counts. For this proof, we assume an undirected network where nodes have two states: pro-vaccine (V_s - score 1), and anti-vaccine (N - score 0), also given in (2.3). Then we have that

$$x_k^n = x_k \quad \forall n \in \mathbb{N}, \quad (\text{A.1})$$

where x_k represents the score of the k th agent. ω is the adjacency matrix of the network, so that $\omega_{j,k} = 1$ if agents j and k are social neighbours, and $\omega_{j,k} = 0$ otherwise. Using the expression for \mathcal{I} given in (2.4) as

$$\mathcal{I} = \frac{N}{W} \cdot \frac{\sum_{j,k} \omega_{jk} (x_j - \bar{x})(x_k - \bar{x})}{\sum_j (x_j - \bar{x})^2}, \quad (\text{A.2})$$

we can derive the expression in (2.5) by first expanding the numerator of (A.2) as

$$\sum_{j,k} \omega_{jk} (x_j - \bar{x})(x_k - \bar{x}) = \sum_{j,k} \omega_{jk} (x_j \cdot x_k - (x_j + x_k) \cdot \bar{x} + \bar{x}^2), \quad (\text{A.3})$$

$$= \sum_{j,k} \omega_{jk} x_j \cdot x_k - \sum_{j,k} \omega_{jk} (x_j + x_k) \cdot \bar{x} + \sum_{j,k} \omega_{jk} \bar{x}^2, \quad (\text{A.4})$$

$$= \underbrace{\sum_{j,k} \omega_{jk} x_j \cdot x_k}_{(\text{I})} - 2 \cdot \bar{x} \cdot \sum_{j,k} \omega_{jk} x_k + \sum_{j,k} \omega_{jk} \bar{x}^2, \quad (\text{A.5})$$

(I) $x_j \cdot x_k$ is nonzero when both $x_j, x_k \neq 0$, so then $j, k \in V_s$; $\sum_{j,k} \omega_{jk} x_j \cdot x_k = 2 \cdot [V_s, V_s]$.

$$= 2 \cdot [V_s, V_s] - 2 \cdot \bar{x} \cdot \underbrace{\sum_{j,k} \omega_{jk} x_k}_{(\text{II})} + \sum_{j,k} \omega_{jk} \bar{x}^2, \quad (\text{A.6})$$

(II) $\omega_{j,k} x_k$ has value when $k \in V_s$ and agent j is any neighbour, so that $\sum_{j,k} \omega_{jk} x_k$ considers the number of neighbours of each vaccinator, and $\sum_{j,k} \omega_{jk} x_k = 2 \cdot [V_s, V_s] + [N, V_s]$.

$$= 2 \cdot [V_s, V_s] - 2 \cdot \bar{x} \left(2 \cdot [V_s, V_s] + [N, V_s] \right) + \bar{x}^2 \cdot \underbrace{\sum_{j,k} \omega_{jk}}_{(\text{III})}, \quad (\text{A.7})$$

(III) $\sum_{j,k} \omega_{jk}$ counts the number of adjacencies between all agent pairs j and k , so that $\sum_{j,k} \omega_{jk} = W$, where W gives twice the number of unique undirected edges of the network.

$$= 2 \cdot [V_s, V_s] - 2 \cdot \bar{x} \left(2 \cdot [V_s, V_s] + [N, V_s] \right) + W \cdot \bar{x}^2. \quad (\text{A.8})$$

For the denominator of (A.2), we have

$$\sum_j (x_j - \bar{x})^2 = \sum_j x_j^2 - 2x_j \cdot \bar{x} + \bar{x}^2, \quad (\text{A.9})$$

$$= \sum_j x_j^2 - 2 \sum_j x_j \cdot \bar{x} + \sum_j \bar{x}^2, \quad (\text{A.10})$$

$$= \sum_j x_j - 2 \cdot \bar{x} \underbrace{\sum_j x_j}_{(\text{IV})} + \bar{x}^2 \underbrace{\sum_j 1}_{(\text{V})}, \quad (\text{A.11})$$

(IV) x_j has value only when $j \in V_s$, so that $\sum_j x_j = [V_s]$ gives the number of pro-vaccine agents.

(V) $\sum_j 1$ sums every node in the network, so that we get the number of nodes in the network;

$$\sum_j 1 = N.$$

$$= [V_s] - 2\bar{x} \cdot [V_s] + N\bar{x}^2, \quad (\text{A.12})$$

$$= (1 - 2\bar{x}) \cdot [V_s] + N\bar{x}^2, \quad (\text{A.13})$$

$$= \gamma \cdot [V_s] + N \cdot \bar{x}^2, \quad (\text{A.14})$$

where

$$\gamma = 1 - 2\bar{x}. \quad (\text{A.15})$$

Therefore, the full expression of Moran's I is written

$$\mathcal{I} = \frac{N}{W} \cdot \frac{2 \cdot [V_s, V_s] - 2\bar{x} \cdot (2 \cdot [V_s, V_s] + [N, V_s]) + W \cdot \bar{x}^2}{\gamma \cdot [V_s] + N \cdot \bar{x}^2}, \quad (\text{A.16})$$

which can be seen as a linear combination of join counts

$$\mathcal{I} = \frac{N}{(\gamma \cdot [V_s] + N \cdot \bar{x}^2) \cdot W} \left((2 - 4\bar{x}) \cdot [V_s, V_s] - 2\bar{x} \cdot [N, V_s] + W \cdot \bar{x}^2 \right). \quad (\text{A.17})$$

Using the expression for the Geary's C statistic \mathcal{C} given in (2.6) as

$$\mathcal{C} = \frac{N-1}{W} \frac{\overbrace{\sum_{j,k} \omega_{jk} (x_j - x_k)^2}^{(\text{VI})}}{\sum_j (x_j - \bar{x})^2}, \quad (\text{A.18})$$

we can identify the numerator (VI) with the expression given for $[N, V_s]$ in (2.2), so that

$$\sum_{j,k} \omega_{jk} (x_j - x_k)^2 = 2 \cdot [N, V_s]. \quad (\text{A.19})$$

Since the denominator of (VI) is identical to that of (A.2), the entire expression for \mathcal{C} can be written as

$$\mathcal{C} = \frac{N-1}{W} \frac{2 \cdot [N, V_s]}{\gamma \cdot [V_s] + N \cdot \bar{x}^2}. \quad (\text{A.20})$$

A.2 Comparisons of models V1, V2 and V3

A.2.1 Model parameters

We use three disparate models V1, V2 and V3 to investigate these EWS, with their respective parameter values and ranges shown in Tab. A.1. Section 2.3 (Results) gives the results for model V2, and here we compare the results of the other models V1 and V3.

Parameter	Interpretation	V1	V2	V3
N	number of agents	10000	40000	562500
ι	case importation (proportion of susceptibles)	2.5×10^{-4}	1×10^{-5}	2.5×10^{-5}
p	infectivity	0.2, 0.8	0.2, 0.8	0.8
ξ_*	random sentiment change	$\xi_1 = 1 \times 10^{-4}$	$10^{-4} \leq \xi_{52} \leq 0.1$	
$\langle d_n \rangle$	mean neighbourhood size	30	50	
σ	strength of social norm	$0 \leq \sigma \leq 3$	$0 \leq \sigma \leq 6$	
κ	perceived risk of adverse vaccine effects	$-1 \leq \kappa \leq 1$	$-2 \leq \kappa \leq 6$	
\mathcal{T}	length of each realisation	$10000 \leq \mathcal{T} \leq 40000$ and until equilibrium	11000	
α	initial pro-vaccine proportion	0.05, 0.95		
ℓ	duration of illness	2 weeks (time steps)		
μ	birth/death rate	2.4×10^{-4}		

Table A.1: A table of the baseline parameter values used for each simulation. N represents the number of agents in the simulation, ξ_1 represents the probability of random sentiment switch per time step, and ξ_{52} represents the probability of randomly switching sentiment once per year (52 time steps).

Models V1 and V3 were parametrised in the way outlined in Section 2.3 (Results): $\kappa = 0$, $\sigma = 0$ and $\alpha = 0.05$ give the result $\langle R \rangle < 0.05$ at equilibrium (defined in Section 1.2). The uniform birth/death rate $\mu = 2.5 \times 10^{-4}$ was chosen to give an average life expectancy of 80 years, and a duration of illness $\ell = 2$ was chosen since the length of the infectious period. Different sizes N were chosen to allow for a sensitivity analysis, given that the three models present major differences in infectivity (p), mean neighbourhood size ($\langle d_n \rangle$), hence the number of effective contacts per week)

and noise (ξ_1 and ξ_{52}), as can be seen in Tab. A.1.

All the transitions shown occur for almost identical values of the social norm σ , suggesting that variance of the noise parameter ξ_{52} has minimal effect over the dynamics of V3. The dynamics of model V2 was shown to be sensitive to slight change in the perceived vaccine risk κ when $\sigma = 0$ (Fig. 2.2). Models V1 (Fig. A.2) and V3 (Fig. A.1) feature similar instabilities to model V2.

For model V3, Fig. A.1 shows that a slight increase in κ from 0 (Fig. A.1(b,e)) to 0.01 (Fig. A.1(a,d)) pushes the system to an anti-vaccine consensus when $\xi_{52} = 1 \times 10^{-4}$, while a decrease to -0.01 (Fig. A.1(c,f)) gives a shift to the opposite pro-vaccine consensus. The similarity in behaviour (trends in the time series) between the panels of Figs. A.1(a-f) ($\xi_{52} = 1 \times 10^{-4}$), A.1(g-l) ($\xi_{52} = 1 \times 10^{-3}$) and A.1(m-r) ($\xi_{52} = 1 \times 10^{-2}$) show that increasing the noise present in the model doesn't affect this sensitivity. Similar observations hold for model V1 (Fig. A.2) for both infectivities $p = 0.2$ (A.2(a-f)) and $p = 0.8$ (A.2(g-l)), where slight changes to the perceived vaccine risk κ result in either anti-vaccine (A.2(a,g)) or pro-vaccine (A.2(c,i)) consensus. One difference between models V2 and V3 is the phenomena of Fig. A.2(d,j), where the increase in p from 0.2 to 0.8 has negatively affected the physical vaccination rate; increasing the infection rate from $p = 0.2$ to $p = 0.8$ takes $[V_p]$ from 0.85 (A.1d) to 0.5 (A.1j).

We stated in Section 2.3 (Results) that the parameter ranges $\kappa \in [-1, 1]$, $\sigma \in [0, 3]$ sufficiently captured transitions in both dynamics as well as the behaviours of the EWS; this is shown for model V2 in Fig. A.3, for models V1 and V3 in Fig. A.4.

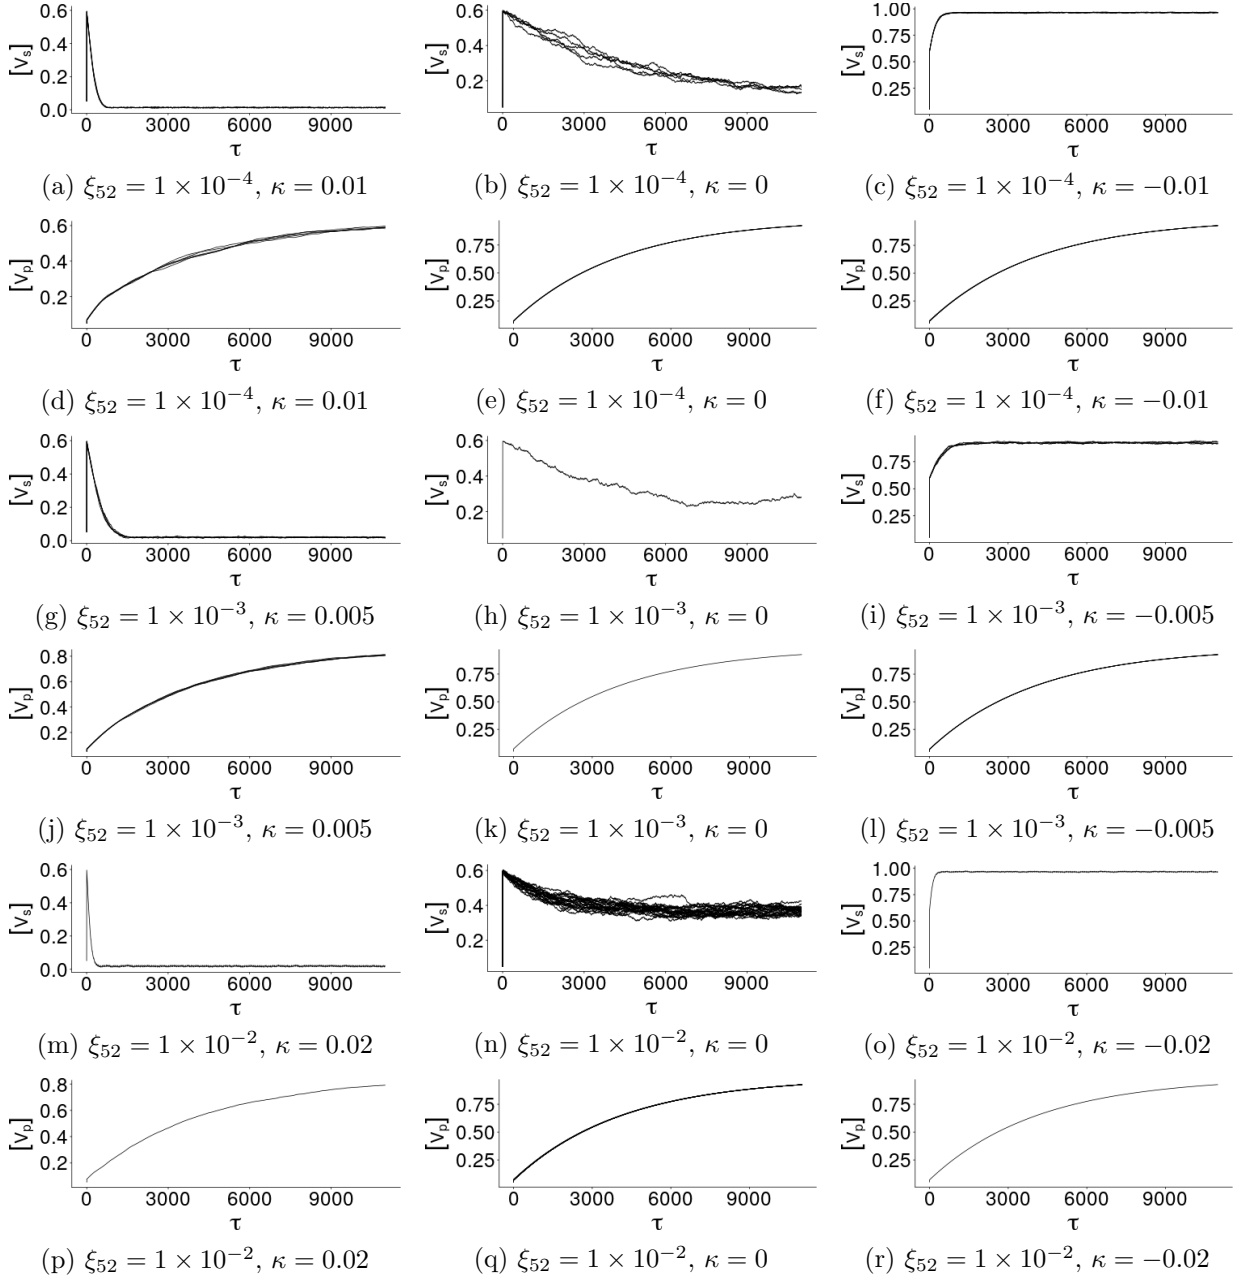


Figure A.1: With no social norm ($\sigma = 0$), varying noise parameter ξ_{52} does not affect model V3 when the perceived vaccine risk κ is close to zero. These time series demonstrate the sensitivity of the social dynamics of model V3 to small changes in κ for various values of ξ_{52} . $[V_p]$ represents the number of vaccinated agents in each time step, and $[V_s]$ gives the number of pro-vaccine agents.

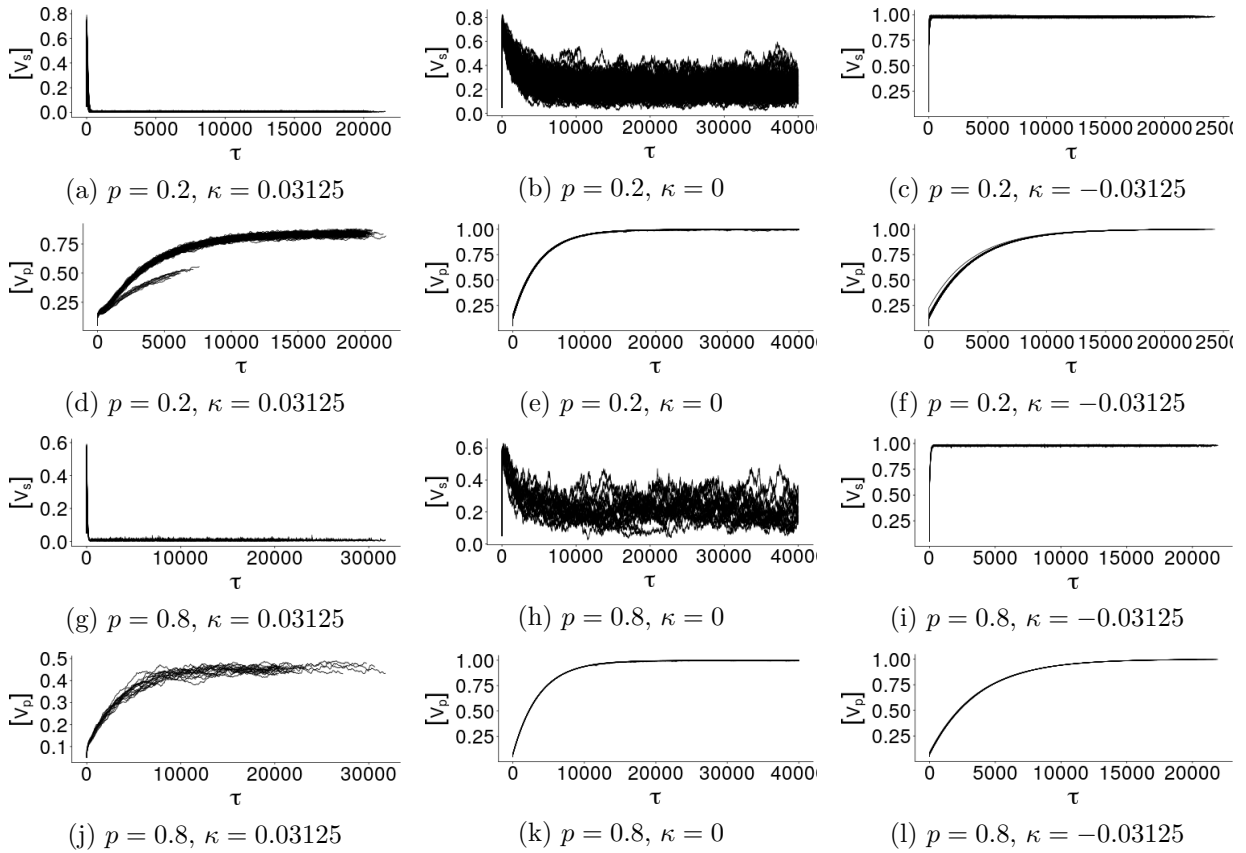


Figure A.2: In the absence of a social norm ($\sigma = 0$), increasing the infectivity of the disease does not alter the sensitivity of the social dynamics of model V1 (similar to models V2 and V3), but does change the vaccination rate when κ becomes positive. $[V_p]$ represents the number of vaccinated agents in each time step, and $[V_s]$ gives the number of pro-vaccine agents. Each panel presents time series from 20 realisations of the Trends of the EWS' equilibrium parameter combination.




Figure A.3: Contour plots of the subregion $(\kappa, \sigma) \in [-1, 0.2] \times [0, 2.4]$ of the parameter space of model V2, showing correspondence between the values of $\langle V_s \rangle$ and $\langle V_p \rangle$, and the proposed EWS. (a) $\langle V_s \rangle$. (b) $\langle V_p \rangle$. (c) \mathcal{M} . (d) $\langle N, N \rangle$. (e) $\langle N, V_s \rangle$. (f) $\langle V_s, V_s \rangle$. (g) $\langle \mathcal{I} \rangle$. (h) $\langle \mathcal{C} \rangle$.




Figure A.4: Contour plots of $\langle V_p \rangle$ and $\langle V_s \rangle$ for the models V1 (a-b) and V3 (c-j) showing that the investigated parameter regions capture the transitions K_s and K_p . (a-b) model V1, $p = 0.8$. (c-d) $\xi_{52} = 0.1$. (e-f) $\xi_{52} = 0.01$. (g-h) $\xi_{52} = 0.001$. (i-j) $\xi_{52} = 0.0001$.

A.2.2 Intertransition distance $K_p - K_s$

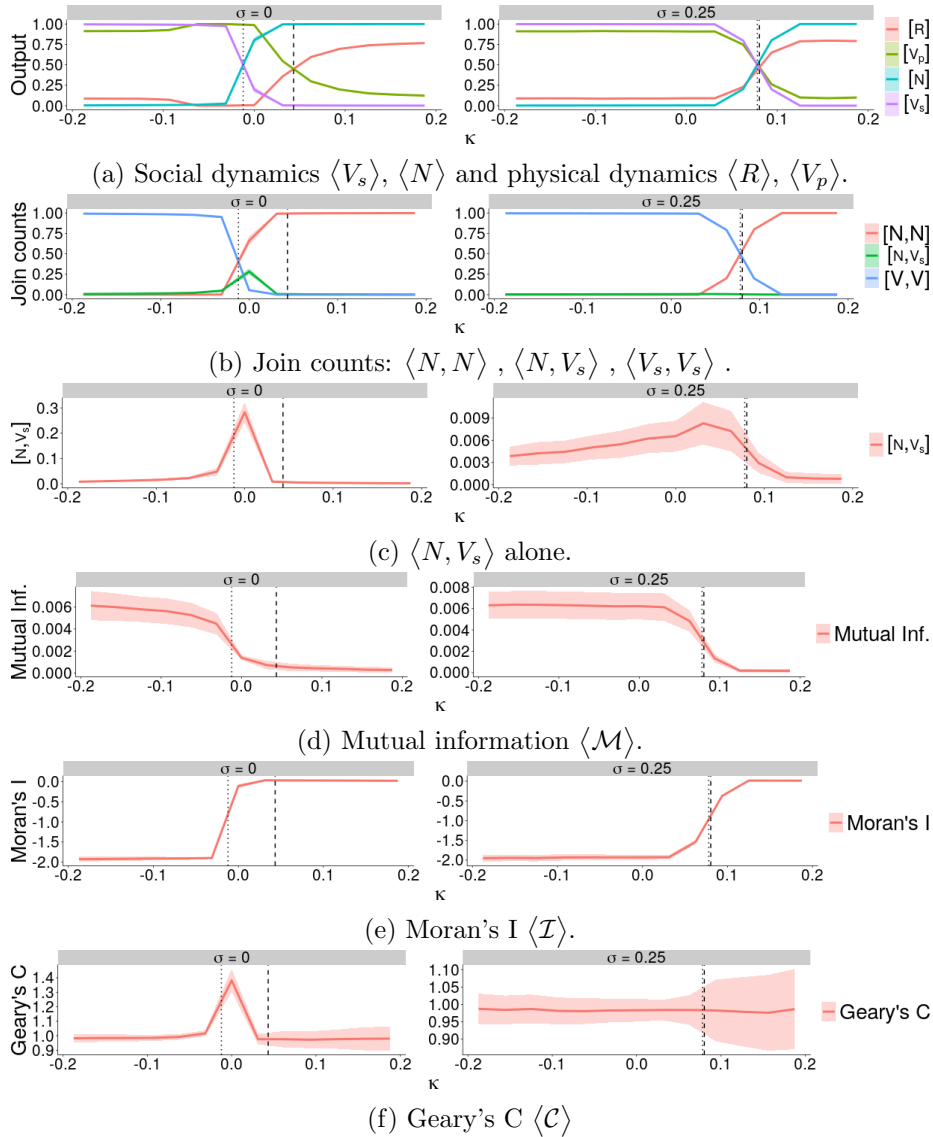


Figure A.5: Trends of the EWS' equilibrium values in model V1 with infectivity $p = 0.2$ approaching the transitions of the social and physical dynamics K_s and K_p (marked by the first and second black lines) respectively, demonstrating the signals given by each tool with respect to the perceived vaccine cost κ . The intervals in each panel represent one standard deviation of the mean equilibrium value in each stochastic realisation of the model.

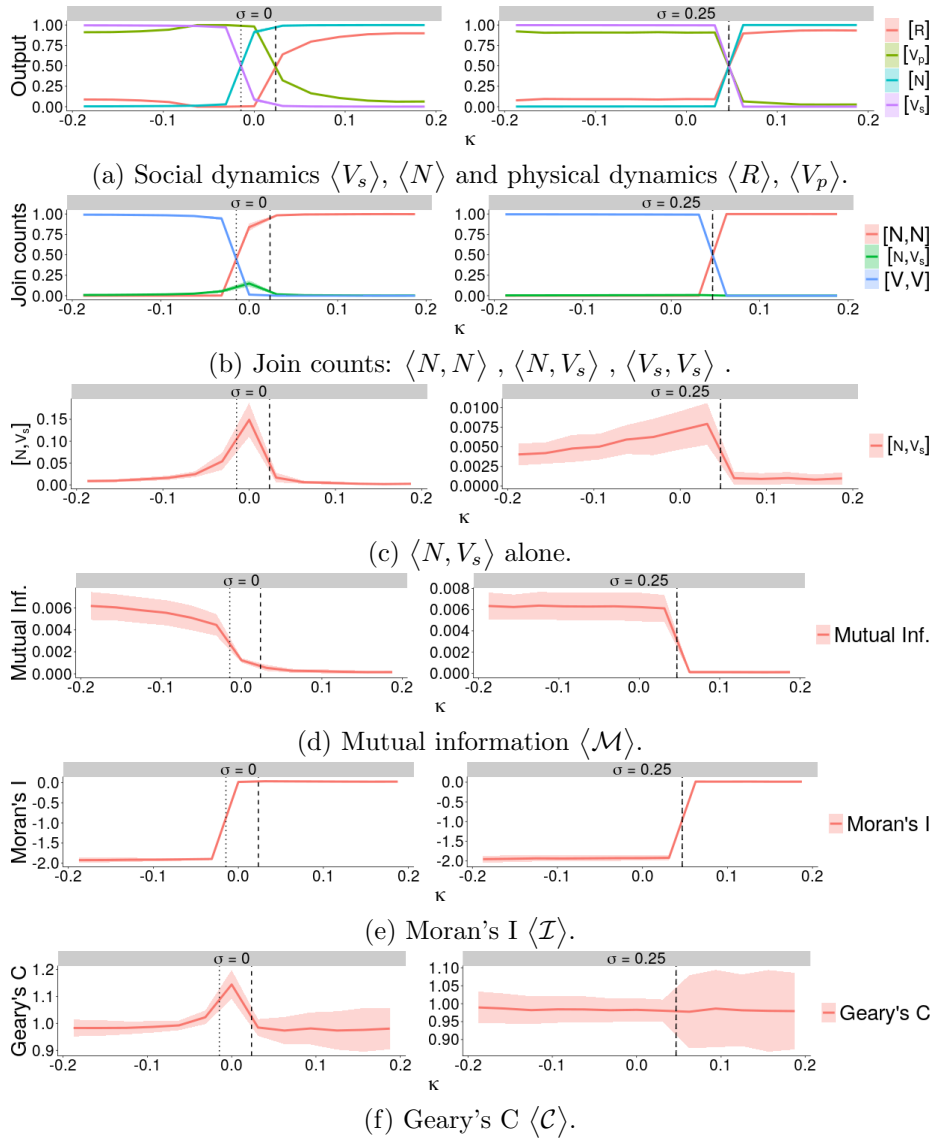


Figure A.6: Demonstration of the trends of the EWS of model V1 with infectivity $p = 0.8$ approaching the transitions of the social and physical dynamics K_s and K_p (marked by the first and second black lines) respectively, demonstrating the signals given by each tool with respect to the perceived vaccine cost κ . The intervals in each panel represent one standard deviation of the mean equilibrium value in each stochastic realisation of the model.

Models V1 and V3 demonstrate a vanishing intertransition distance similar to model V2 in Fig. A.7, where the distance between the dotted and dashed vertical lines (representing K_s and K_p) in Fig. 2.3a, left decreases as the social norm gets stronger in Fig. 2.3a, right; this shrinking distance between the two vertical lines can also be seen in Figs. A.5 and A.6 by comparing the first and second columns. The independence of these trends in the intertransition distance to the difference in the three models was expected; they follow immediately from previous discussions in which few substantive differences were seen. Since this is then a property of the dynamics themselves, any discussion about the interpretation and validity of the EWS will be largely identical among the three models.

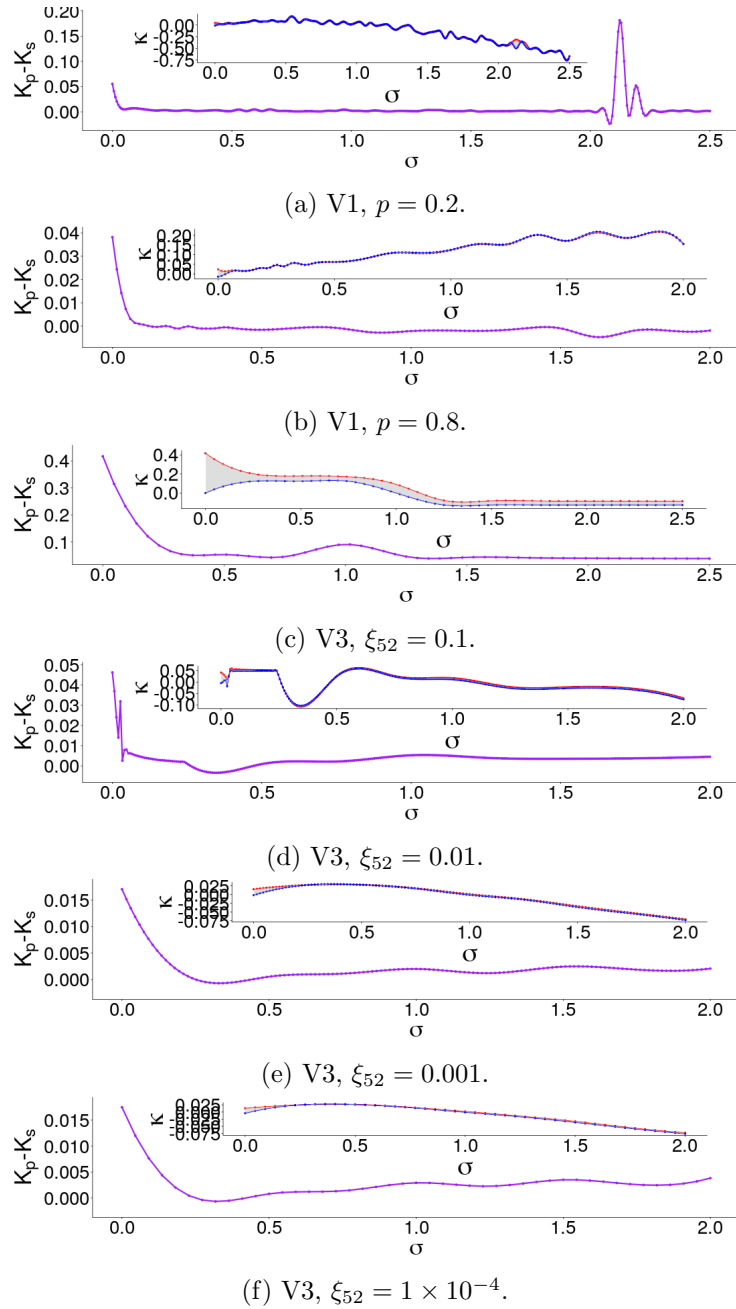


Figure A.7: Demonstration of the vanishing intertransition distance $K_p - K_s$ (purple) for all models, with inset graphs showing the estimated locations of K_p (insets, blue) and K_s (insets, red).

A.2.3 Multiple definable transitions K_s and K_p

A difficulty with the interpretation of Fig. 2.5a (and all panels of Fig. A.7) is the case where there are multiple social (K_s) and physical (K_p) transitions. The definitions of K_s and K_p given in Section 2.3 (Results) rely on finding the *earliest* κ value where $\langle V_s \rangle \sim \langle N \rangle$ and $\langle V_p \rangle \sim \langle R \rangle$ respectively. For model V2, this problem occurs in the range $\sigma \geq 1$ (Fig. A.8). Of particular interest are the trends in the proposed EWS; for V2 the behaviour of $\langle \mathcal{M} \rangle$ (mutual information) in this range is shown in column (B) of figure Fig. A.8. There is still a change in trend approaching the transitions K_s and K_p (which occur almost simultaneously), as shown in Fig. 2.5a ($K_p \sim K_s$ when $\sigma > 1$); since both the warnings from the EWS and the definitions of K_s and K_p depend on the lowest κ values, the EWS still give early warning signals despite this.

Figure A.10m details the number of transitions in the physical ($\#K_p$, red) and social ($\#K_s$, blue) dynamics per value of the social norm σ ; the height of each bar represents the total number of transitions for the corresponding value of the social norm σ , and the length of each coloured portion of the bar represents the number of transitions of that type. For example, there are 23 definable transitions when $\sigma = 1.875$; 15 are social transitions (the length of the red portion of that bar), and the remaining 8 transitions occur in the physical dynamics ($23 - 15 = 8$). For model V2, multiple transition ($\#K_* > 1$) occurs in the region $\sigma > 1$. However, our description of a transition also counts occurrences where the two opposing curves merely touch each other (as best as can be determined from an ensemble of stochastic realisations, rather than the result of some rigorous analysis), thereby possibly inflating the value of the number of *meaningful* transitions producing changes in the vaccination rate or aggregate sentiment.

A similar observation can be made with model V1 in the range $\sigma > 0.5$ (Fig. A.9); the trends of the dissimilar join count $\langle N, V_s \rangle$ (column B) still demonstrate statistically significant change before either transition despite the irregularity of the trends shown (as compared to those occurring in the region $\sigma \geq 0.75$). Fig. A.9(i-j) show a marked increase in the intertransition distance (the gap between the two horizontal lines), also shown by the sudden change in trend shown in Fig. A.7a in the range $2 < \sigma < 2.125$. This represents a region of the parameter space of model V1 where the physical transition K_p lagged the social transition K_s far more than expected due to the occurrence of a social transition *without an accompanying physical transition*. This assertion is supported by Fig. A.9m, where the bar chart detailing the numbers of transitions shows that there are many more social transitions than physical transitions ($\#K_s > \#K_p$). Previously, even though there were multiple transitions K_s and K_p , the first social transition was always a reliable predictor of the first physical transition; here, there are multiple social transitions before the first physical transition is seen; this presents a different opportunity for the misjudgement of the proximity of a vaccine crisis, since this may lead to a large lead distance in parameter regions presenting otherwise smaller lead distances. For example, if we expect a vaccine crisis to follow soon after a shift in aggregate sentiment, then a detected social transition may lead to time-sensitive preparations for a vaccine crisis that occurs far later than expected.

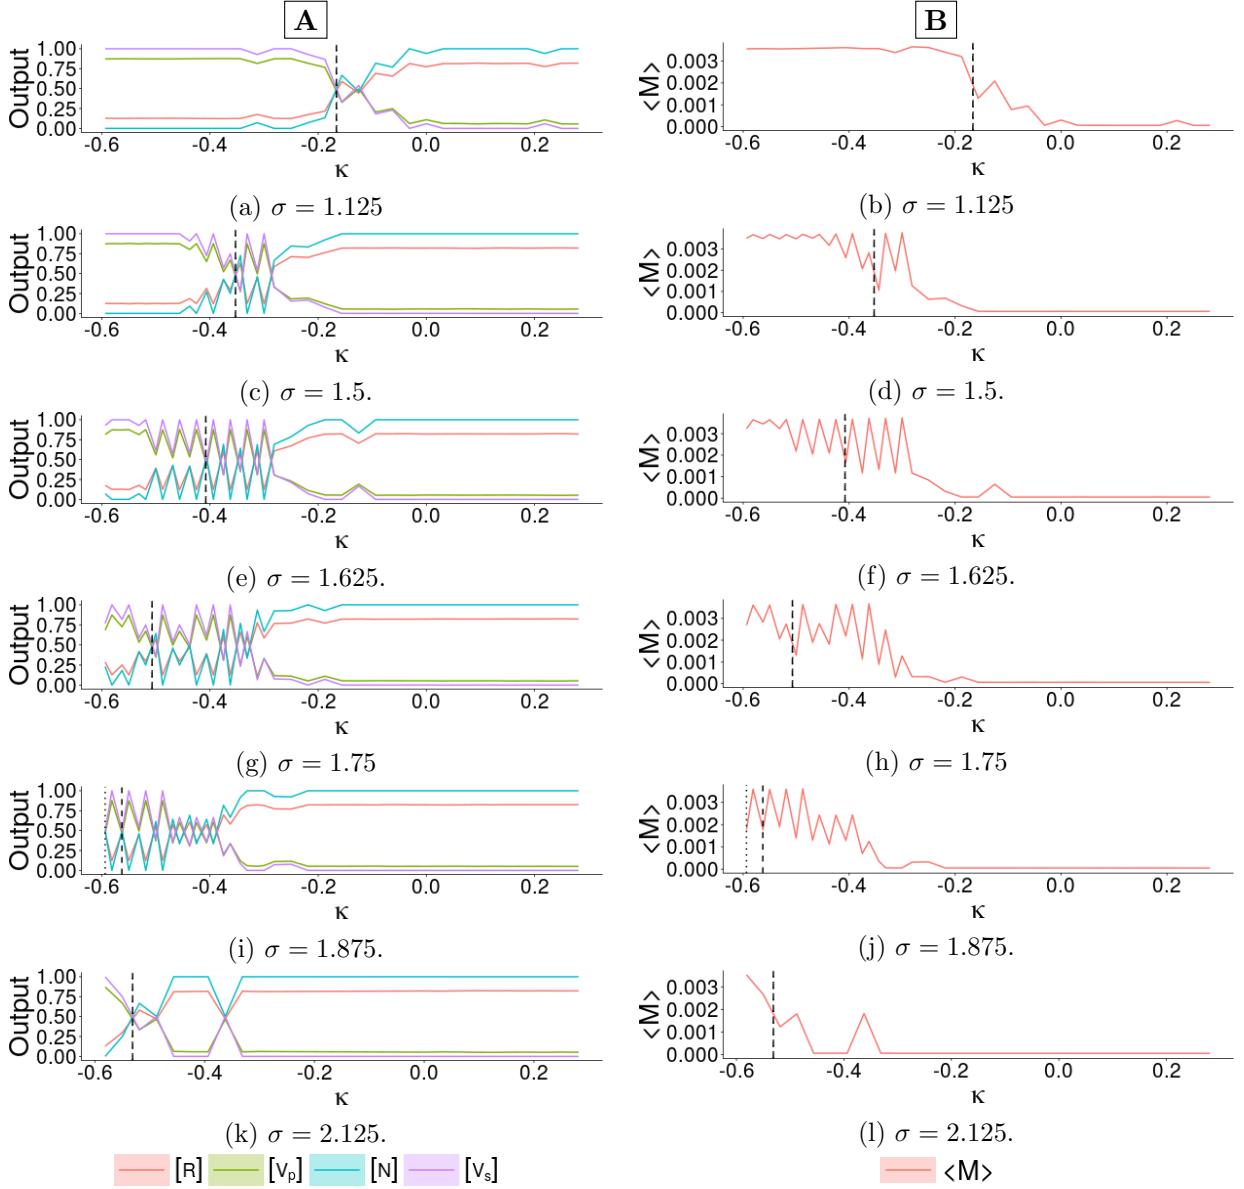


Figure A.8: For model V2, multiple potential transitions K_s and K_p can be identified for $\sigma > 1$, introducing the opportunity for false warnings unless the earliest possible transition is chosen (in either dynamic), as was done in the main text. This region features erratic behaviour of $\langle \mathcal{M} \rangle$ in regions featuring multiple physical and social transitions. The vertical dotted and dashed vertical lines give the locations of K_s and K_p respectively.

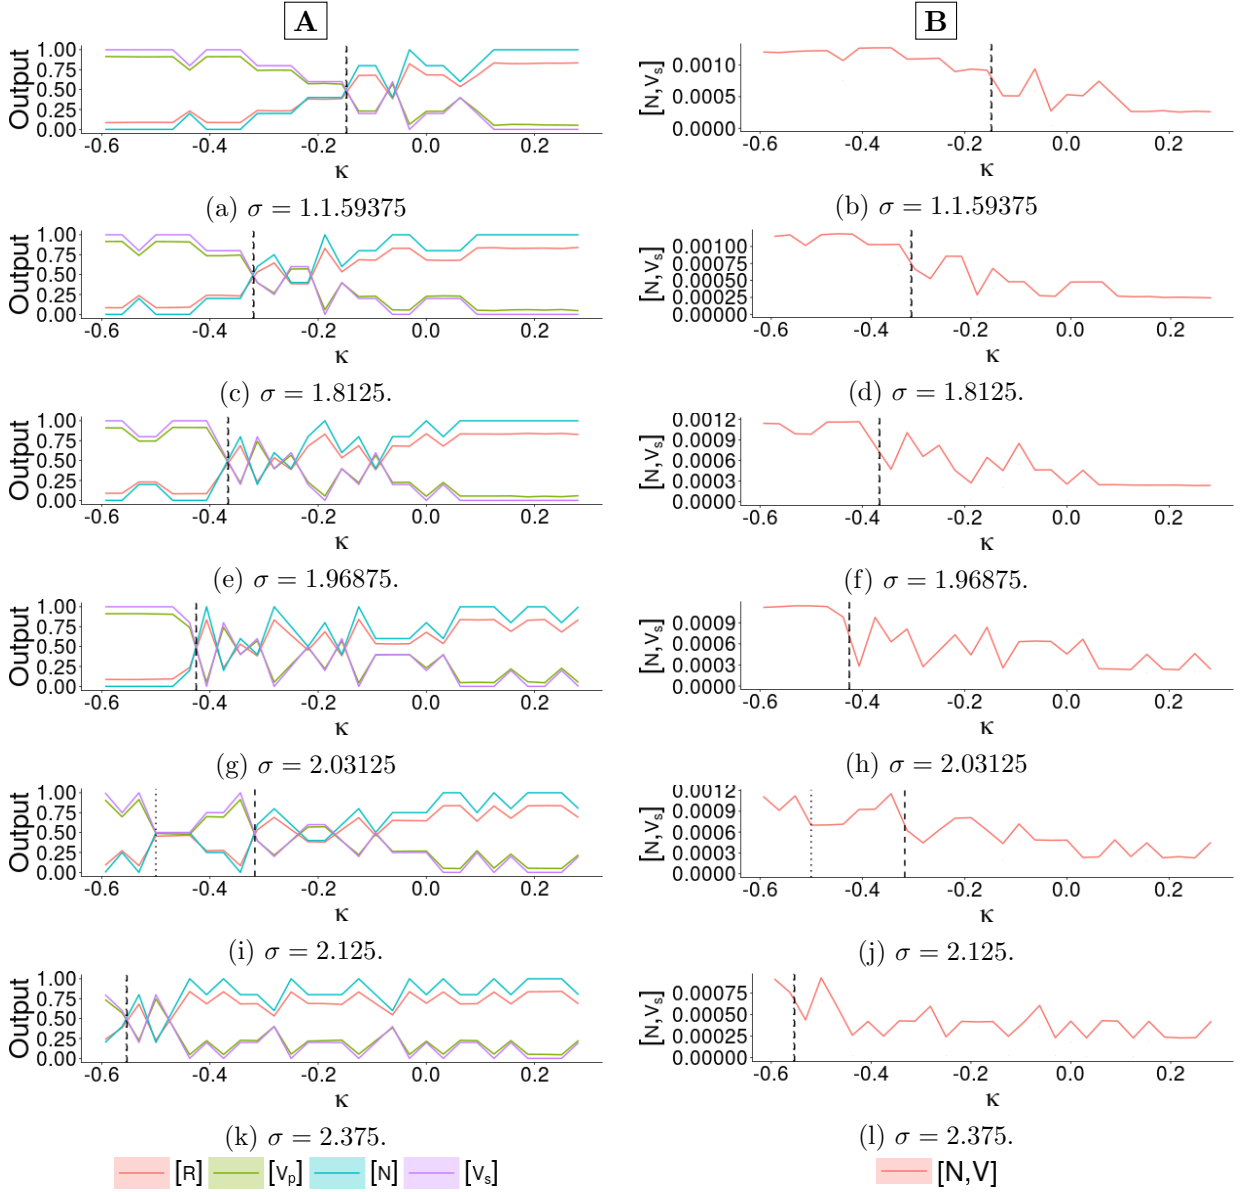




Figure A.9: Similar to model V2 (Fig. A.8), multiple social ($\#K_s > 1$) and physical ($\#K_p > 1$) transitions occur in realisations of model V1 ($p = 0.2$) at $\sigma \geq 0.75$, again providing opportunities for false warnings. Erratic behaviour of $\langle N, V_s \rangle$ occurs in regions featuring multiple physical and social transitions. (i-j) present an instance of multiple social transitions occurring before the first physical transition, presenting an opportunity for the misjudgement of the proximity of a vaccine crisis. The dotted and dashed vertical lines give the location of K_s and K_p respectively.



(a) Model V1.



(b) Model V2.

Figure A.10: Bar chart showing the number of physical ($\#K_p$, red) and social ($\#K_s$, blue) transitions with respect to the value of σ .

In sum, all the values of K_s and K_p estimated in this study may not represent the earliest definable transition in the dynamics, but rather the first definable transition *within the range* $-1 \leq \kappa \leq 1$. As will be shown in App. A.2, different models (specifically different values of N) show a wider range of transitions of the values K_s and K_p .

A.2.4 Change point testing and warnings of the EWS

Since 15 – 20 realisations were run for each unique tuple of parameter values, averaged time series were made by finding the mean of all time series at each time step t . For each value of σ , a corresponding κ -series was formed by assembling the means of the last 500 time steps of each averaged time series for each value of κ . Per value of σ , each change point test was applied to the κ -series as follows: the series for each EWS was ordered in increasing values of κ , and the change point test was applied to successively longer segments of the sequence (of length at least 3, starting from the first value $\kappa = -1$) until a predicted change point B_* was returned with a certainty $p < 0.05$. Four change point tests were used: the Buishand test[90], the Pettitt test[347, 441], the Lanzante test[319] (used in the main text), and the Standard Normal Homogeneity test[14, 15] (hereon referred to as the *SNHT*).

This method can be justified by the various panels of Fig. A.11, which show how the warnings of each EWS depend on the length of the κ -series of values the test was applied to. For example, the purple line and points in panel Fig. A.11l show the warning given by the $\langle N, N \rangle$ count under the SNHT test for model V2. Starting (in all instances) from $\langle V_s \rangle_{\kappa=-1}$, applying the SNHT to a series of length 4 gives a statistically insignificant result ($p \approx 0.8$); there is a corresponding purple filled point at $n = 4$ on the x -axis, and the same explanation holds for $4 \leq n \leq 7$. For $n = 8$ on the x -axis (representing a sequence of length 8), we see the first statistically significant result under the SNHT ($p = 0.02$), represented by the start of the purple line.

As progressively longer sequences of $\langle V_s \rangle_{\kappa \geq -1}$ are used ($n \uparrow$), the SNHT continues to give predictions with p values under the 0.05 threshold of significance (as evidenced by the continuation of the line in Fig. A.11l, rather than the sudden appearance of filled points), though $\text{SNHT}_\sigma\{\langle \Psi \rangle\}$ itself is increasing. Interestingly, as $n \rightarrow 30$, $\text{SNHT}_\sigma\{\langle \Psi \rangle\}$ approaches the value of K_s , and the other lead distances ($\text{BR}_\sigma\{\Psi\}$ in red, $\text{Lan}_\sigma\{\Psi\}$ in blue and $\text{Pet}_\sigma\{\Psi\}$ in green) also show the same increasing trend in the final moments before the arrival of K_s ; all panels of Fig. A.11 end at $n = 32$, since this is when the end point of the κ -series used for each change point test reaches or passes K_s (there is little use here for predictions of K_s given *after* the transition has occurred).

For model V2, comparing the panels of **column C** of Fig. A.11 shows that the $\langle N, N \rangle$ and $\langle N, V_s \rangle$ counts (Figs. A.11l and o respectively) give the earliest significant warnings under the SNHT (as seen from the early start of the solid purple lines representing $\text{SNHT}_\sigma\{\langle \Psi \rangle\}$). Conversely, the latest warnings come from $\langle V_s, V_s \rangle$ (Fig. A.11r) under all four change point tests, since the earliest significant result occurs at $n = 18$. This seems to be due to the relative value of the joint counts; $\langle V_s, V_s \rangle = \mathcal{O}(N^2)$ while $\langle N, N \rangle$ and $\langle N, V_s \rangle$ are both small before K_s , so that the establishment of a trend on $\langle N, N \rangle$ and $\langle N, V_s \rangle$ will be “more noticeable” than that for $\langle V_s, V_s \rangle$.

While the signals of the joint counts $\langle N, N \rangle$, $\langle N, V_s \rangle$ and $\langle V_s, V_s \rangle$ increase under all change point tests as K_s is approached (Fig. A.11(j-r)), $\langle \mathcal{C} \rangle$, $\langle \mathcal{I} \rangle$ and $\langle \mathcal{M} \rangle$ remain relatively stable as more terms in the κ -series are added to the test (Fig. A.11(a-i)). Our interpretation of this difference

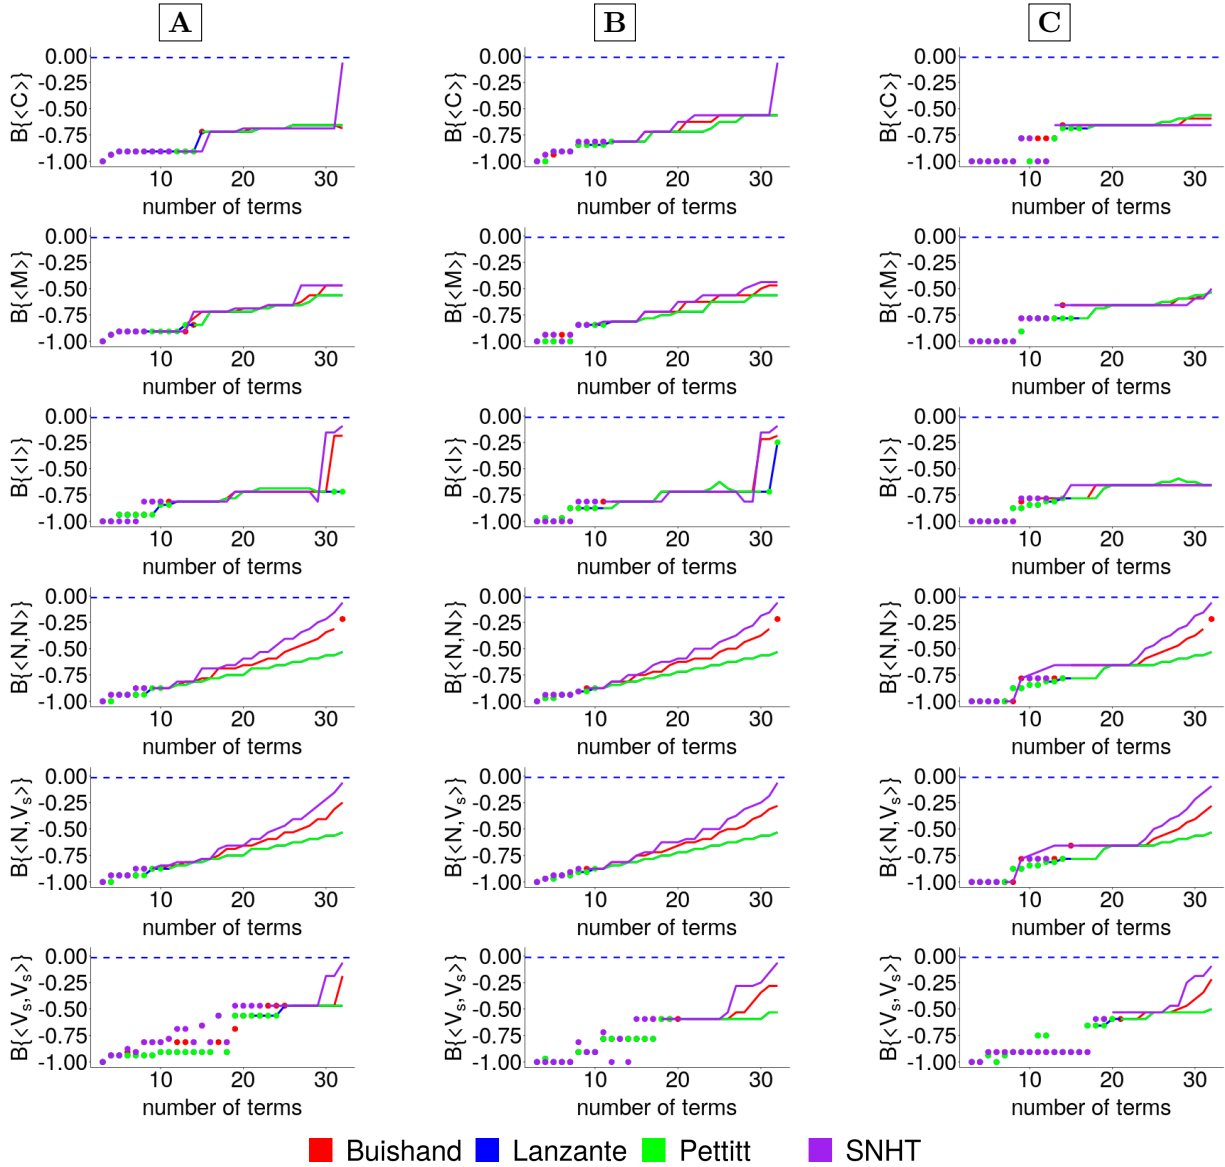


Figure A.11: Trends of the predicted warnings of the EWS with respect to the length of the κ -series used. Within each panel, solid lines are estimated change points with $p < 0.05$ (significant), while insignificant estimates ($p \geq 0.05$) are represented by filled points. K_s is indicated by a blue dashed horizontal line in each panel. (column A) model V1 with $p = 0.2$. (column B) model V1 with $p = 0.8$. (column C) model V2. (a-c) the application of the change point tests to the κ -series for Geary's C $\langle C \rangle$. (d-f) Mutual information $\langle M \rangle$. (g-i) Moran's I $\langle I \rangle$. (j-l) Anti-vaccine similar join count $\langle N, N \rangle$. (m-o) Dissimilar join count $\langle N, V_s \rangle$. (p-r) Pro-vaccine similar join count $\langle V_s, V_s \rangle$.

is neither positive or negative, since a stable prediction inspires trust in the result, whereas the establishment of an increasing trend in the warnings of the join counts seems a better indicator of the proximity of K_s than the estimate given by the test (this is a case where the initial value and trends of the result are both useful). Since this increase in the warnings of $\langle N, N \rangle$, $\langle N, V_s \rangle$ and $\langle V_s, V_s \rangle$ (Fig. A.11(j-r)) persists under all the change point tests used, this trend can be seen as a property of the κ -series of the join counts *and hence the join counts themselves as EWS* rather than a spurious observation. A similar analysis holds for model V1 in columns (A) and (B) of Fig. A.11, where again $\langle N, N \rangle$ and $\langle N, V_s \rangle$ give the earliest warnings of the approach of K_s (Fig. A.11(j,m)) for $p = 0.2$ and (A.11(k,n)) for $p = 0.8$.

Measurements of the lead distance for transitions K_s under the Lanzante test are shown in Figs. 2.5(c-d); here, we show the lead distances given by other change point tests used on the EWS for model V2. Comparing the panels of Figs. A.12, A.13 and A.14, all EWS retain their validity under the various tests, except for the consistent failure of the EWS at higher values of σ for both models (V1 and V2) previously attributed to insufficient range of κ as σ increased. Of all the EWS, the performance of $\langle V_s, V_s \rangle$ and $\langle \mathcal{C} \rangle$ vary the most with respect to the test used for both V1 and V2. For instance, under the SNHT (Fig. A.12g), $\langle \mathcal{C} \rangle$ is the only failure of all the EWS, with a negative lead distance in the range $1.75 \leq \sigma \leq 2.5$ (i.e. the warning occurs post-transition). Also of note is that for model V2, the Lanzante test (Fig. A.12(c-d)) not only gives the highest maximum lead distances for all the EWS, but also yields a consistently larger average lead distance than the other change point tests. It appears that the Lanzante test provides the largest lead distances in the range $\sigma < 2$ for V2.

Similar observations cannot be made for model V1, however; important differences in Figs. A.13 and A.14 are that $\langle V_s, V_s \rangle$ is also now inconsistent (along with $\langle \mathcal{C} \rangle$), with earlier failure than in model V2 at $\sigma = 1.21875$ as well as $\sigma > 1.90625$ when $p = 0.2$. This can possibly be attributed to behaviours due to the different parameter values of the models, and or the finer resolution of σ values in model V1 as compared to model V2; values of σ increase in increments of 0.03125 in model V1, whereas the finest increment of σ is 0.125 in model V2. Similar to model V2, $\langle \mathcal{I} \rangle$ and $\langle N, V_s \rangle$ perform well under all the change tests with relatively high lead distances when $p = 0.2$, but $\langle N, N \rangle$ now also gives a lead distance comparable to the two previous EWS. Finally, there is not as much variation in maximum lead distance per EWS for model V1 as there is for model V2; maximum lead distances for the best EWS ($\langle \mathcal{M} \rangle$, $\langle N, V_s \rangle$ and $\langle V_s, V_s \rangle$) are all around 1.

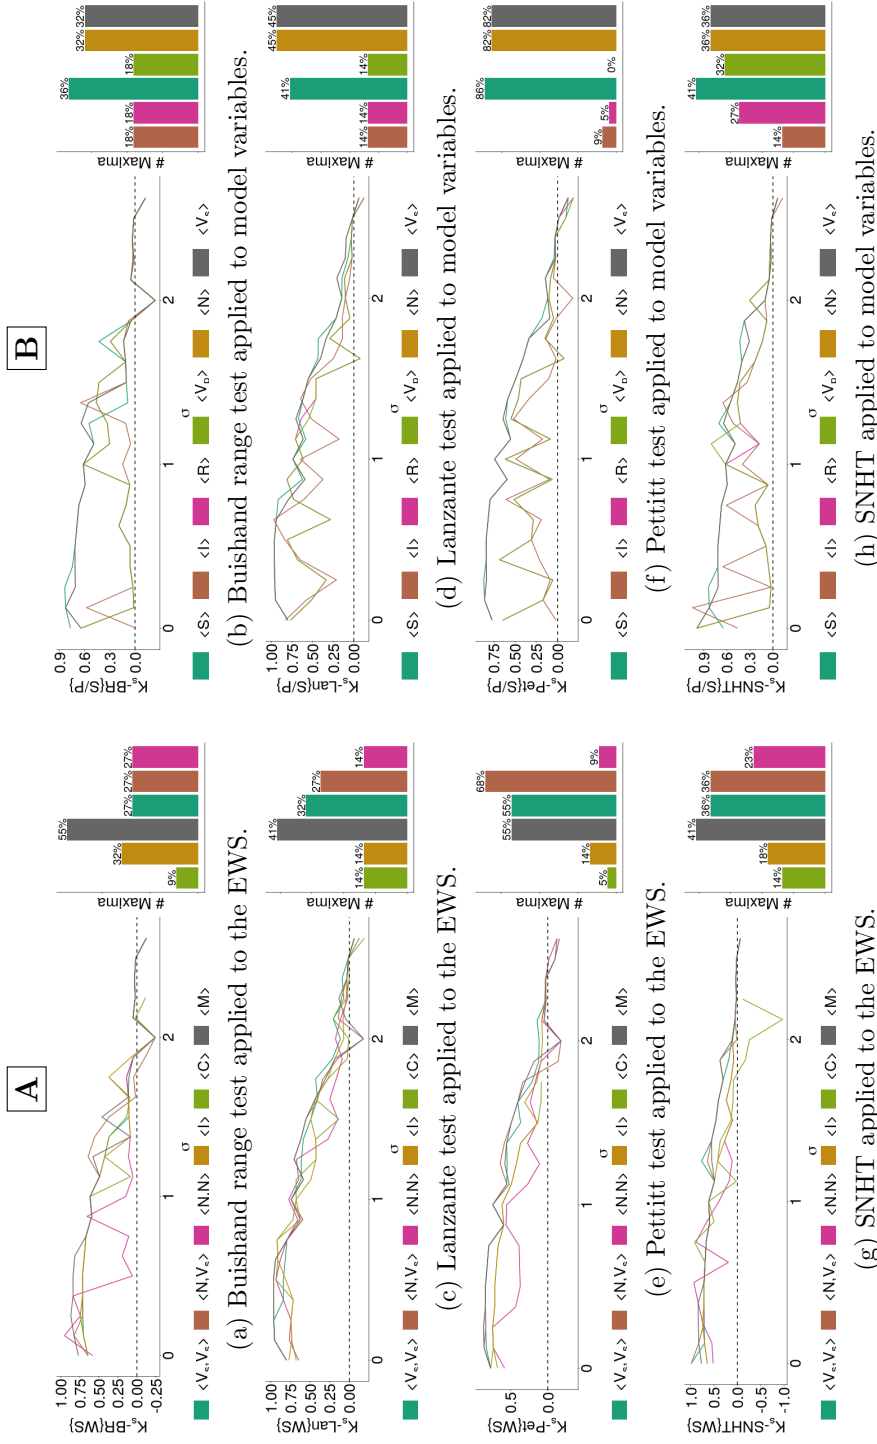


Figure A.12: The lead distances given by different change point detection tests for model V2. The panels of (column A) give the results of the tests applied to the early warning signals (EWS), and (column B) for the model dynamics. (a-b) Buishand range test. (c-d) Lanzante test. (e-f) Pettitt test. (g-h) Standard normal homogeneity test.

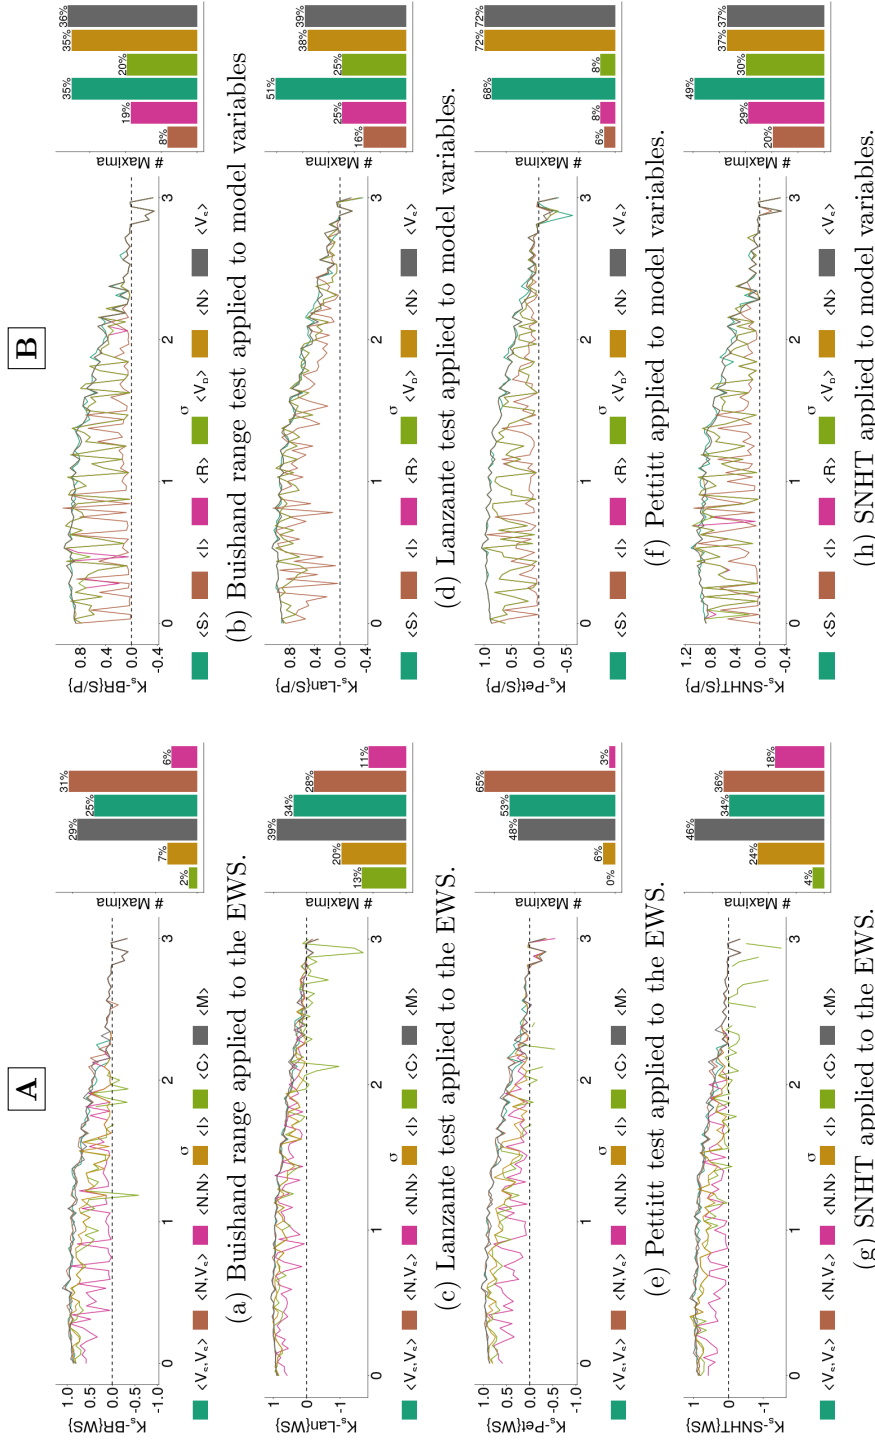


Figure A.13: The lead distances given by different change point detection tests for model V1, with infectivity $p = 0.2$. Similar to Fig. A.12, the panels of column A give the results of the tests applied to the early warning signals (WS), and column B for the model dynamics. (a-b) Buishand range test. (c-d) Lanzante test. (e-f) Pettitt test. (g-h) Standard normal homogeneity test.

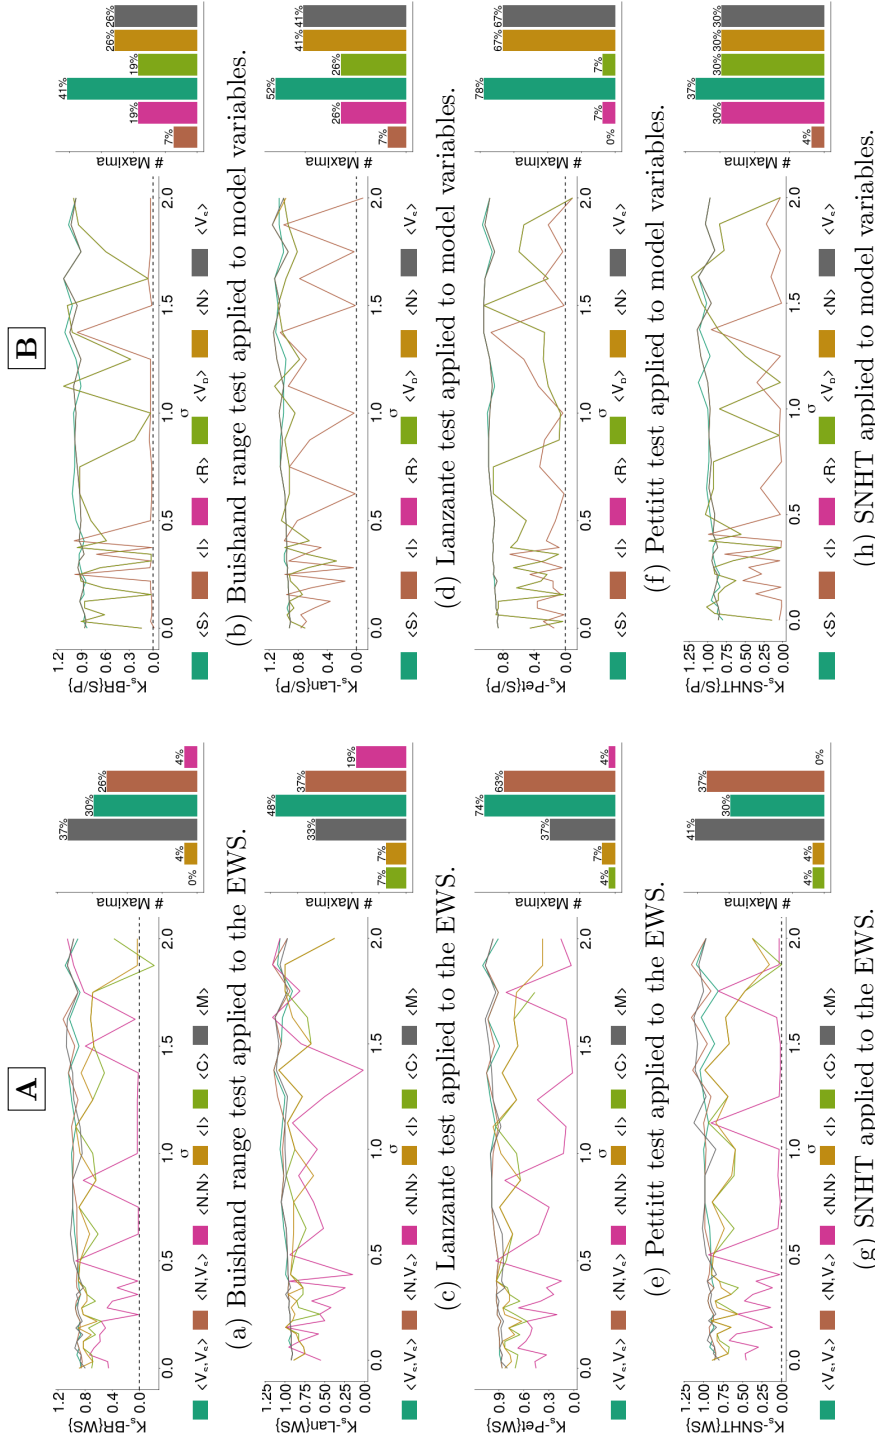


Figure A.14: The lead distances given by different change point detection tests for model V1, with infectivity $p = 0.8$. Similar to Fig. A.13, the panels of column A give the results of the tests applied to the early warning signals (WS), and column B for the model dynamics. (a-b) Buishand range test. (c-d) Lanzante test. (e-f) Pettitt test. (g-h) Standard normal homogeneity test.

A.2.5 Reversibility of the κ -series of the EWS

Figure A.15(column A) plots the skewness γ_1 of each trend with respect to the intertransition distance $K_p - K_s$ (column B) plots the skewness γ_1 against the strength of the social norm σ ; (column B) suggests a relationship between the skewness γ_1 and the social norm σ , though it may not be causative. As an example of skewness in the case of a weaker social norm $\sigma \approx 0$, trends in the mutual information $\langle \mathcal{M} \rangle$ (Fig. A.15b), Moran's I $\langle \mathcal{I} \rangle$ (Fig. A.15d) and Geary's C $\langle \mathcal{C} \rangle$ (Fig. A.15e) are symmetric (i.e., small γ_1) for small σ . As σ increases however, the direction of κ becomes vital to the calculation and interpretation of warning signals; for example, (Fig. A.15j) shows a clear change in the skewness of the dissimilar join count $\langle N, V_s \rangle$ as $\sigma \rightarrow 0.25$. This change is visible moving between Figs. 2.3e ($\sigma = 0$) and 2.3f ($\sigma = 0.25$), where the trend becomes more left-skewed.

Figure A.16 shows the skewness of the κ -series of the EWS with respect to the social norm σ for models V1 with $p = 0.2$ (Fig. A.16(column A)), $p = 0.8$ (Fig. A.16(column B)) and V2 (Fig. A.16(column C)). This quantification of the change of the shape of the κ -series with increasing σ suggests a concern of directionality in the application and interpretation of the EWS; there is no hint that these EWS would be of comparable effectiveness if the model had dynamics such that the derived κ -series were reversed.

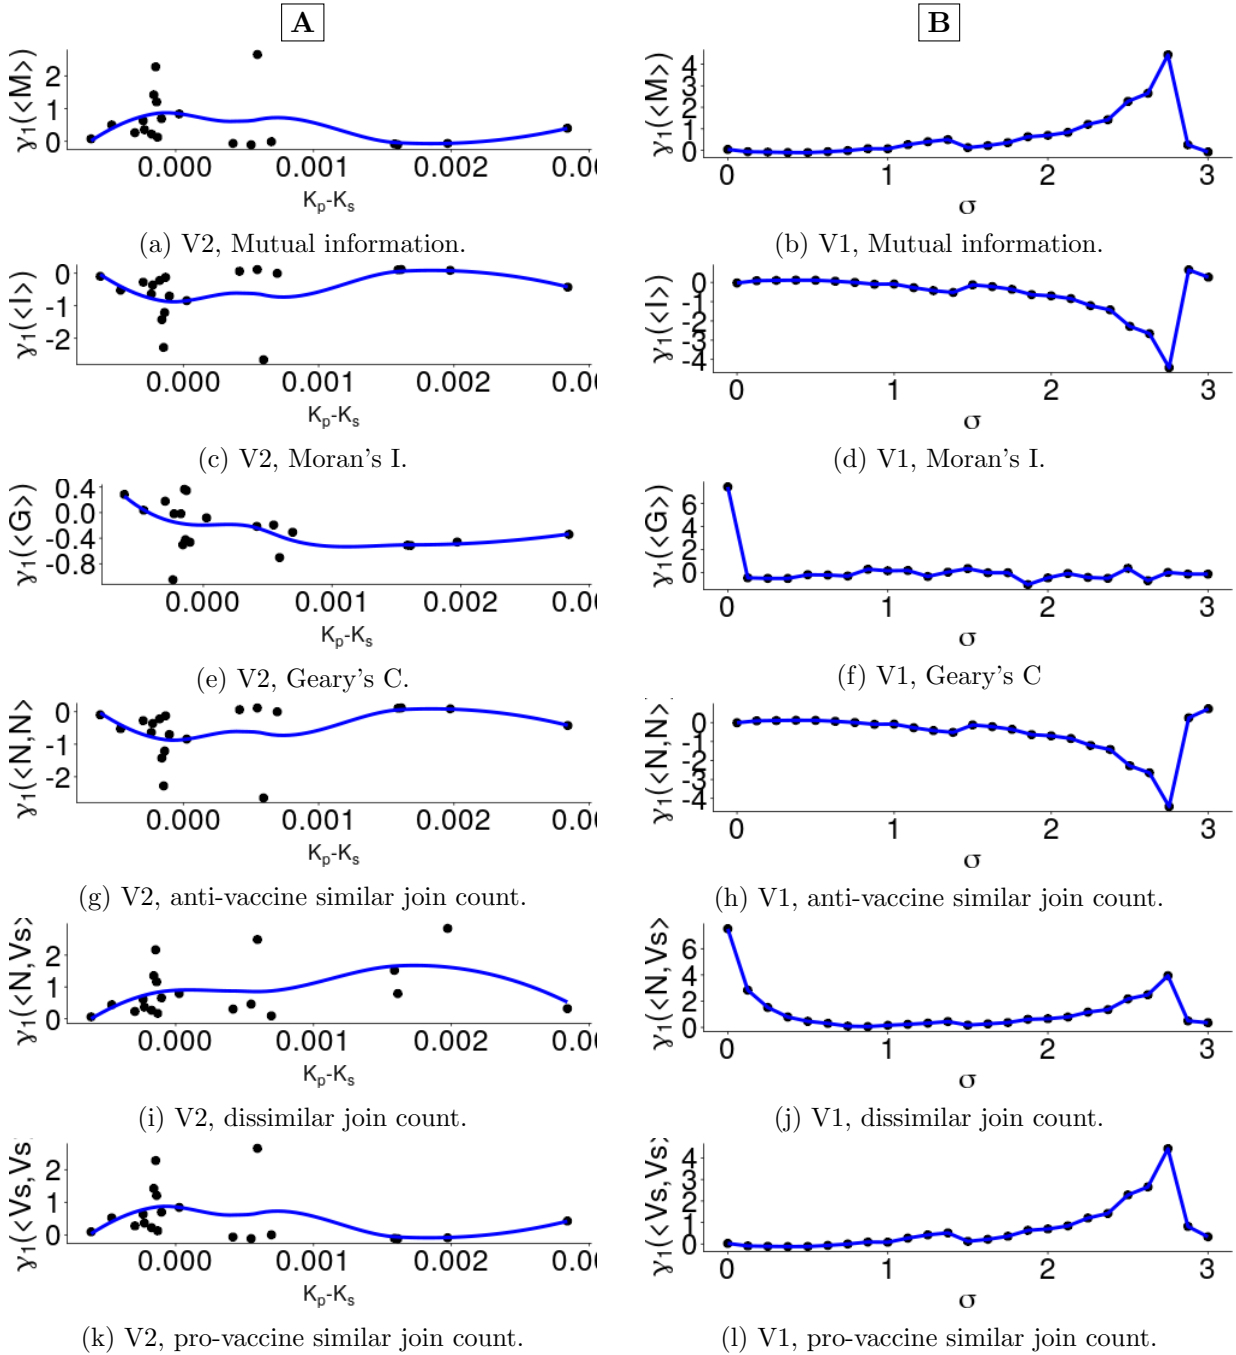


Figure A.15: The skewness of the EWS' trends varies more consistently with the strength of the social norm σ (column B) than with the intertransition distance $K_p - K_s$ (column A) for model V2. κ series such as those shown in Figs. 2.3 , A.8 and A.9 are generally asymmetric about both K_s and K_p .

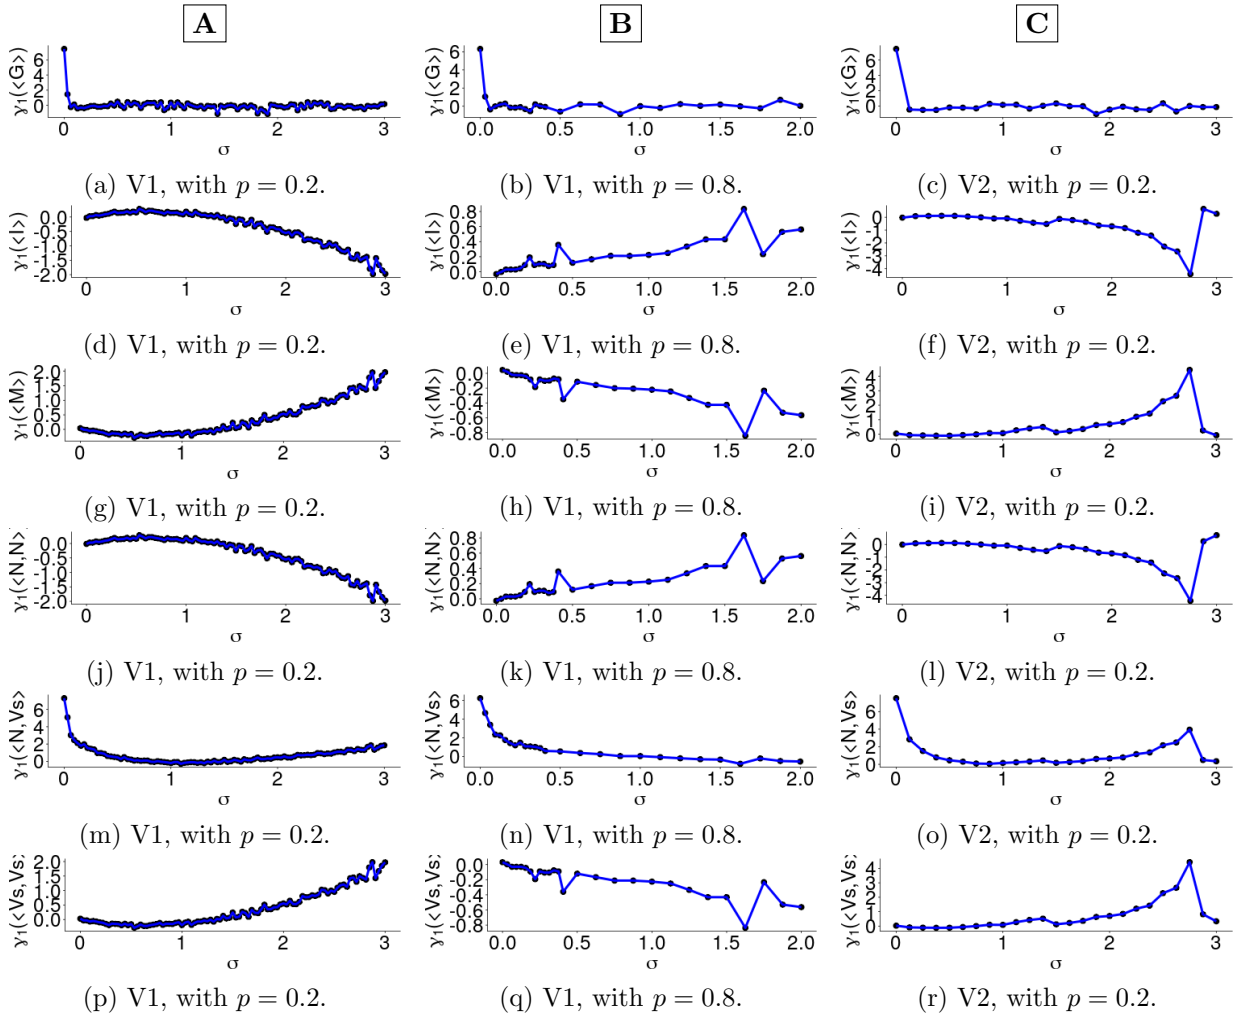


Figure A.16: Trends of the skewness γ_1 of the κ -series with respect to the strength of the social norm σ for models V1 with $p = 0.2$ (column A), $p = 0.8$ (column B) and V2 (column C). (a-c) give the skew of $\langle C \rangle$, (d-f) $\langle \mathcal{I} \rangle$, (g-i) $\langle \mathcal{M} \rangle$, (j-l) $\langle N, N \rangle$, (m-o) $\langle N, V_s \rangle$, (p-r) $\langle V_s, V_s \rangle$.

A.2.6 Further comparisons of the EWS' performance

We use the same definitions and measures of performance χ_{\min}^* (maximin) and χ_{\max}^* (maximax) given in Section 2.3 (Results). Fig. 2.7 shows the results of maximax and maximin comparisons of the performances of the EWS and S/I classes of signals using the Lanzante change point detection test for model V2. Figure A.17 shows this comparison using all various change point detection tests (Buishand, Lanzante, Pettitt and Standard normal homogeneity tests), and Fig. A.18 makes these comparisons for model V1. Points falling in the green-shaded regions represent σ values at which either the best-performing (largest lead distance) EWS outperforms all model variables ($\chi_{\max}^* > \epsilon_{\max}^*$), or where the worst-performing EWS still outperforms the worst-performing variable ($\chi_{\min}^* > \epsilon_{\min}^*$). Points in the red-shaded region represent the reverse.

Both Fig. A.17 and Fig. A.18 show equal performance of the EWS and the model variables in both maximin and maximax comparisons, though all panels of Fig. A.18 together show many more instances of equal maximum lead distances ($|\chi_{\max}^*| < \epsilon_{\max}^*$, red points and curve) among the two classes in model V1 than in V2 Fig. A.17. Also important is the changes in area of the green- and red-shaded regions among the panels of Fig. A.18 and Fig. A.17; for instance, a large green-shaded region (such as in Fig. A.18g) shows that the maximum lead distance of the EWS is much larger than that of the model variables for those σ values for which the EWS perform better.

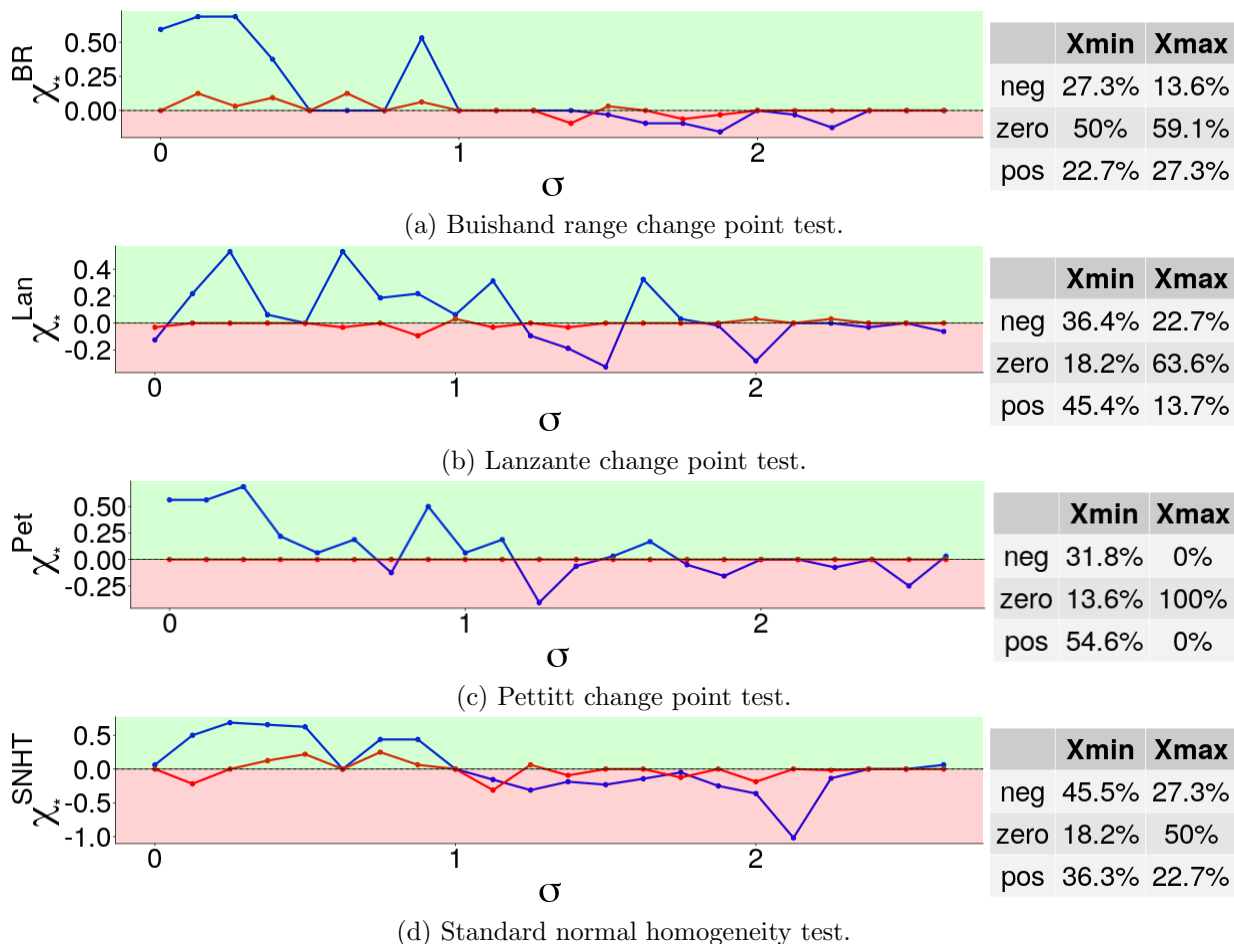


Figure A.17: Comparisons of the performance of EWS (WS) and model dynamics (S/I) for model V2 with respect to four change point detection tests. In each panel, χ_{\min}^* is represented by a blue curve through blue points, and χ_{\max}^* by a red curve through red points. The green-shaded region represents the region where either the worst-performing EWS (that is, giving the smallest lead distance) still outperforms with worst-performing model variable ($\chi_{\min}^* > \varepsilon_{\min}^*$) under the specific change point test, or where the best-performing EWS outperforms the best-performing model variable ($\chi_{\max}^* > \varepsilon_{\max}^*$). The inset table gives the ratios of social norm σ values for which $\chi_{\min}^* < -\varepsilon_{\min}^*$ (neg), $|\chi_{\min}^*| < \varepsilon_{\min}^*$ (zero) and $\chi_{\max}^* > \varepsilon_{\max}^*$ (pos).

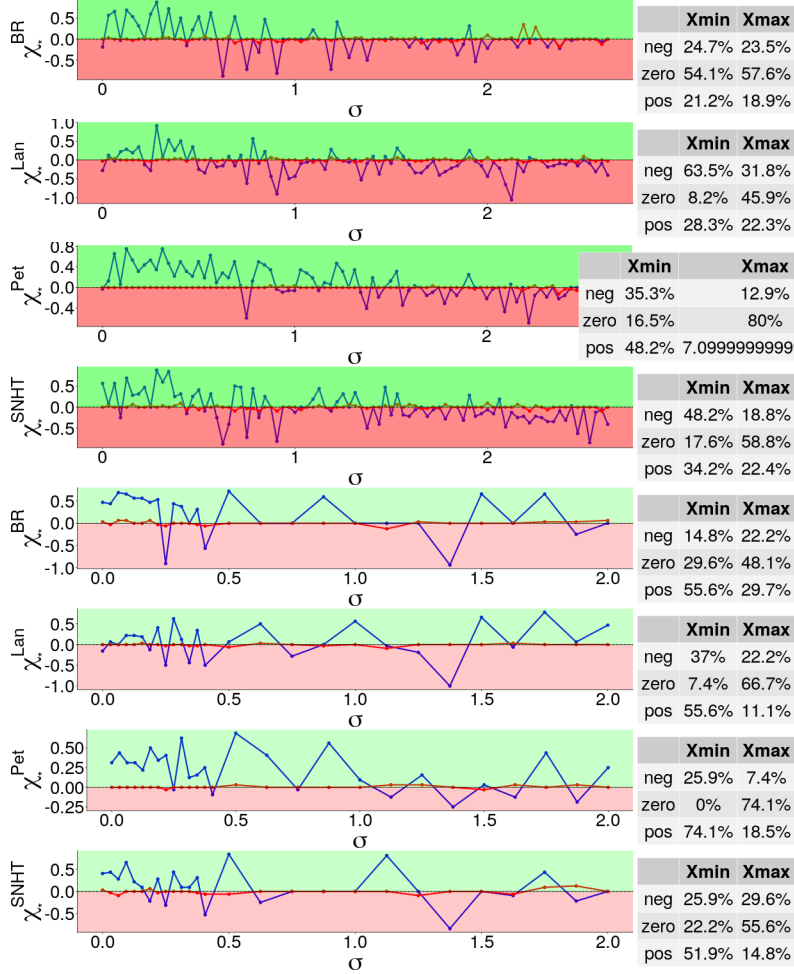


Figure A.18: Similar to Fig. A.17, we show comparisons of the performance of EWS (WS) and model dynamics (S/I) for model V1 with respect to four change point detection tests: (a,e) Buishand range test (BR). (b,f) Lanzante test (Lan). (c,g) Pettitt test (Pet). (d,h) Standard normal homogeneity test (SNHT). (a-d) infectivity $p = 0.2$, (e-h) infectivity $p = 0.8$. In each panel, χ_{\min}^* is represented by a blue curve through blue points, and χ_{\max}^* by a red curve through red points. The green-shaded region represents the region where either the worst-performing EWS (that is, giving the smallest lead distance) still outperforms with worst-performing model variable ($\chi_{\min}^* > \varepsilon_{\min}^*$) under the specific change point test, or where the best-performing EWS outperforms the best-performing model variable. The inset table gives the ratios of social norm σ values for which $\chi_c^* < -\varepsilon_c^*$ (neg), $|\chi_c^*| < \varepsilon_c^*$ (zero) and $\chi_c^* > \varepsilon_c^*$ (pos).

A.2.7 EWS performance per σ value


In Sec. 2.3, the EWS' performance under the Lanzante test (for model V1) was compared by examining the proportion of social norm σ values for which each EWS gave the greatest or smallest lead distances of all EWS. Here we do the same for other change point tests.

Figure A.19 shows the performance of each EWS tested under each of the Lanzante (Fig. A.19b), Pettitt (Fig. A.19c), Buishand range (Fig. A.19d) and the SNHT (Fig. A.19a) change point tests. As can be seen by comparing the different length of the bars corresponding to each EWS, performance varies depending on the change point test used. For instance, mutual information $\langle \mathcal{M} \rangle$ gives the most maxima (largest lead distances, green bar) under the SNHT (Fig. A.19a, 59%) than it does under the Pettitt test (Fig. A.19c, 55%), say; however it gives its smallest proportion of minima (smallest lead distances, red bar) under the Pettitt test than the SNHT. Across all panels of Fig. A.19, the anti-vaccine similar join count $\langle N, V_s \rangle$ shows the worst performance, with its ratio of maxima (green bar) peaking under the SNHT test (Fig. A.19a, 23%), with the least ratio of minima (red bar) occurring under the Pettitt test (Fig. A.19a, 14%).


Figure A.20 compares the proportions of maxima (largest lead distances) and minima (smallest lead distances) for all the EWS of model V1 for different change point tests. As was noted in Fig. A.19, performance varies quite a bit between EWS depending on the test. Generally, the dissimilar join count $\langle N, V_s \rangle$ appears to be the strongest-performing EWS, with consistently high proportions of maxima (green bars) and low proportions of minima (red bars) across all change point tests and infectivities shown in Fig. A.20.

Figure A.21 represents the performance of each EWS at each σ value, and allows us to see whether good or bad performance of each EWS can be restricted to any specific subset of σ values. These grid diagrams are related to the panels of Fig. A.19 as follows; per EWS and change point test, the length of the green bar in the relevant panel of Fig. A.19 gives the proportion of green tiles appearing in the relevant row of a panel in Fig. A.21; the same relationship holds for the proportion of minima (red bars) in the panels of Fig. A.19 and the proportion of red tiles in the relevant row of Fig. A.19. Similar to the observations made in Sec. 2.3, there seems to be no emergent pattern in σ values for which either maxima or minima occur for each EWS. The most appropriate description of each EWS seems to be the overall proportion of maxima and minima.

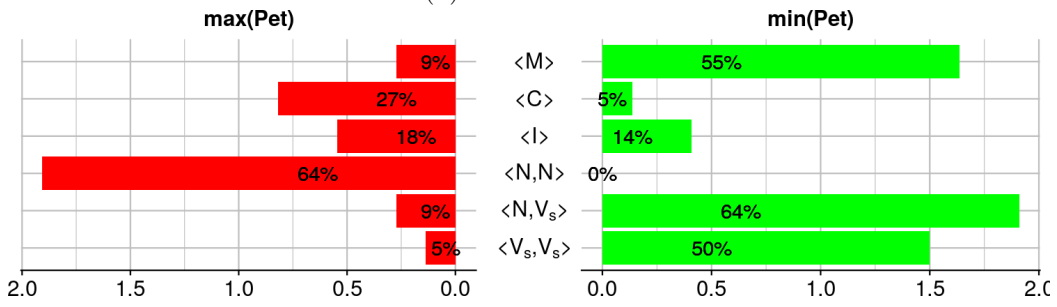
Again, this time for model V1 (Fig. A.22), there is no pattern in each EWS' row of any panel suggesting that any EWS is best used on a specific subset of σ values. The only new information gained here deals with the nature of lead distances given; for example, we can observe that the Geary's C $\langle \mathcal{C} \rangle$ EWS performs badly for $\sigma \geq 1.5$ under the SNHT (Fig. A.22a(left)), Pettitt (Fig. A.22b(left)) and Buishand range (Fig. A.22d(left)) change point tests.



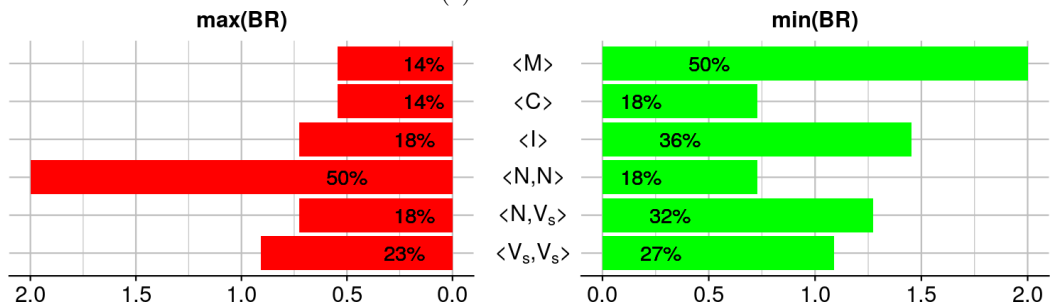
(a) Standard normal homogeneity test.



(b) Lanzante test.



(c) Pettitt test.



(d) Buishand range test.

Figure A.19: Bar charts showing the ratios of σ values for which each EWS gave the largest (green bars) and smallest (red bars) lead distance of all EWS for model V2.

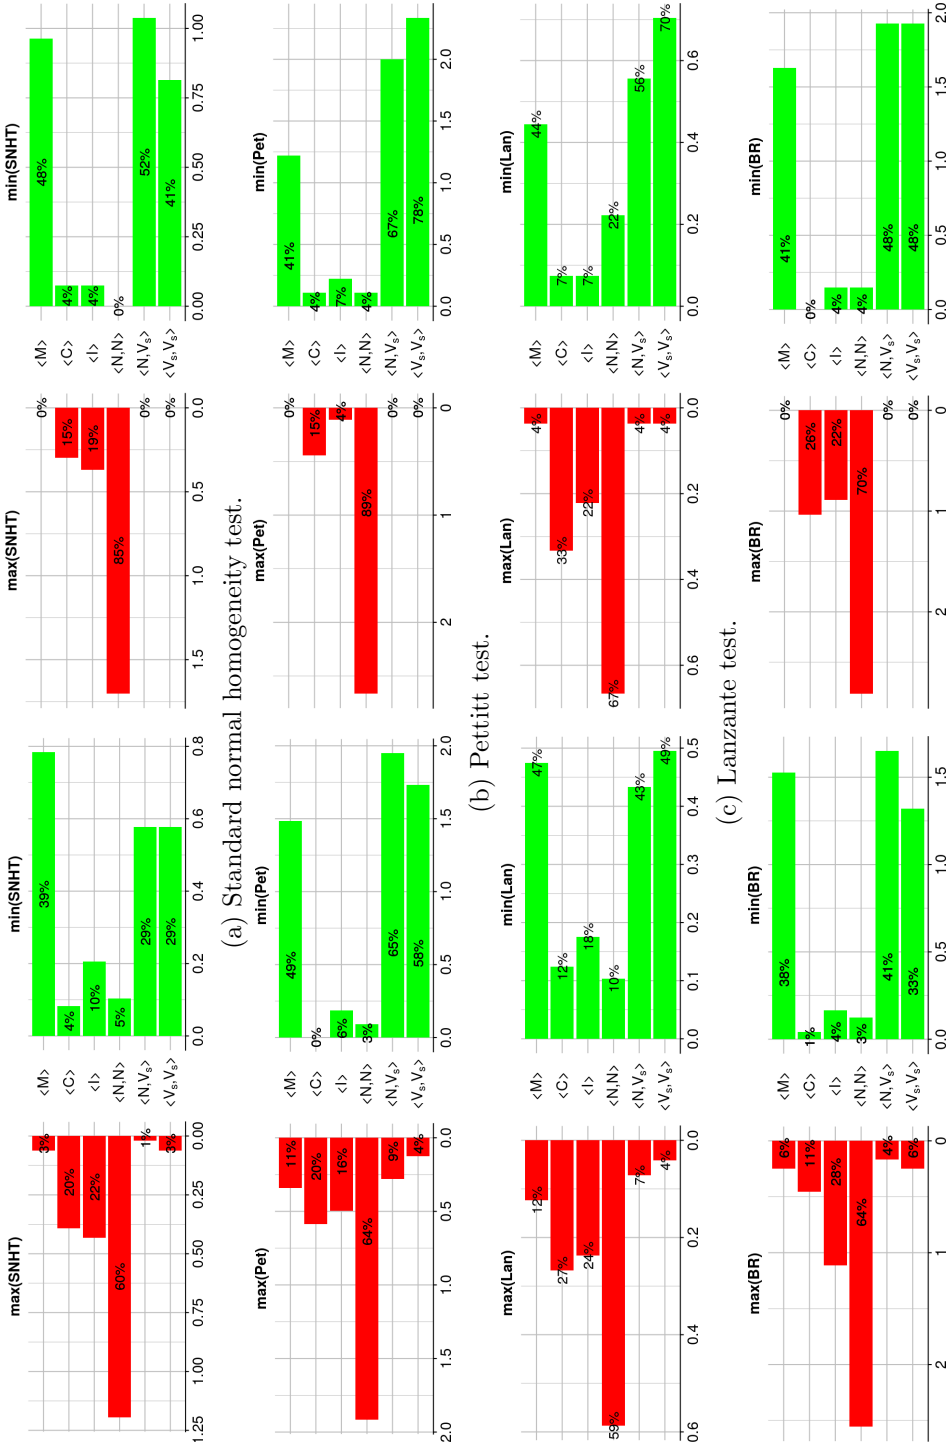





Figure A.20: Bar charts showing the proportions of σ values for which each EWS gives the largest (green bars) and smallest (red bars) lead distances of all EWS for model V1 under each of the Standard normal homogeneity (a), Pettitt (b), Lanzante (c) and Buishand (d) tests. Infectivity is $p = 0.2$ for panels on the left, and $p = 0.8$ for panels on the right.




(a) Standard normal homogeneity test.



(b) Pettitt test.

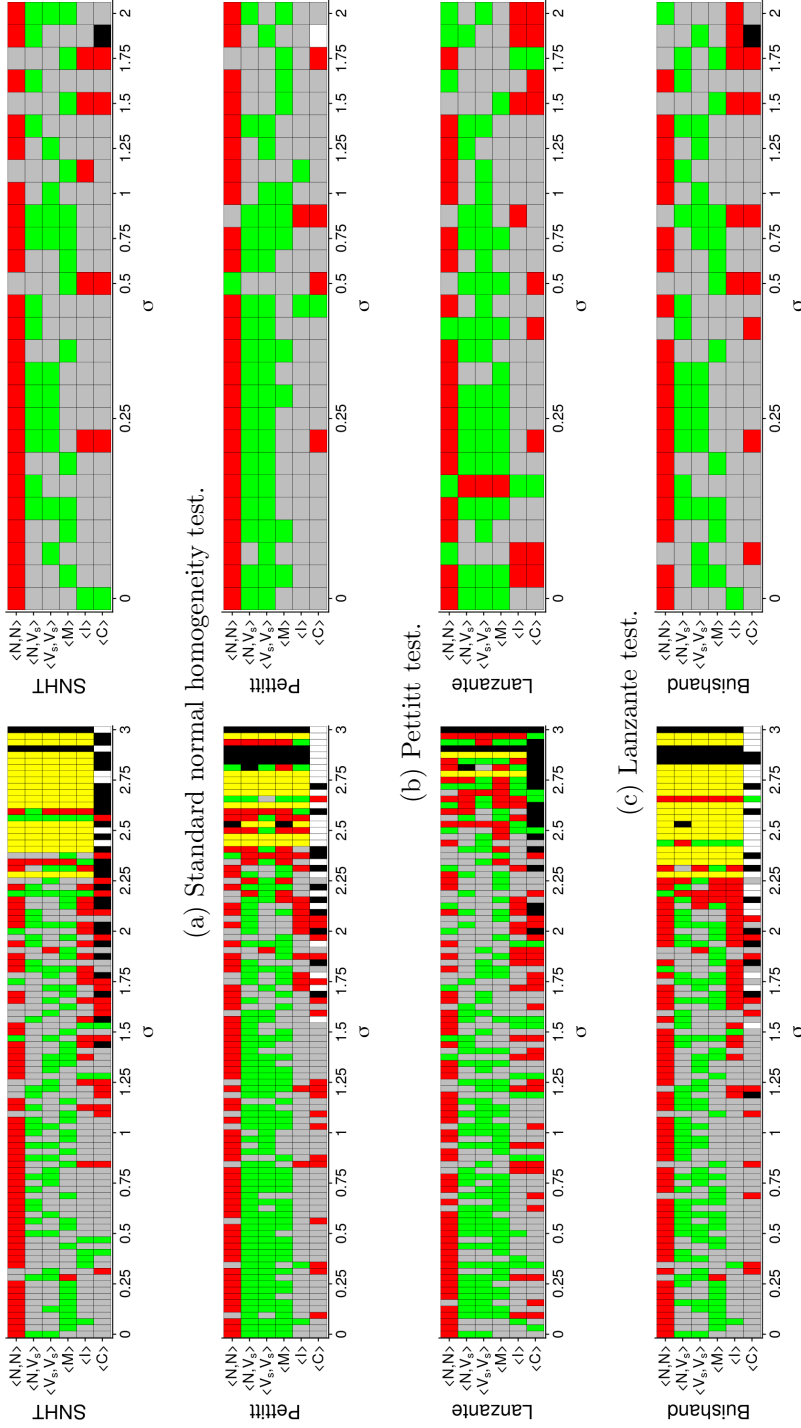


(c) Lanzante test.



(d) Buishand range test.

Figure A.21: Grid plots showing the relative performance of each EWS with respect to the value of the social norm σ for model V2. Green tiles represent a maximum for the EWS (i.e., it gives the largest lead distance of all the EWS), red tiles represent a minimum (that is, the smallest lead distance of all EWS), yellow tiles indicate that all EWS have the same value, grey tiles represent neutral lead distances (neither maximum nor minimum). Black tiles represent failed warnings (negative lead distances), and white tiles represent indeterminate values.



(a) Standard normal homogeneity test.

(b) Pettitt test.

(c) Lanzante test.

(d) Buishand range test.

Figure A.22: Grid plots showing the relative performance of each EWS with respect to the value of the social norm σ for model V1. Green tiles represent a maximum for the EWS (i.e., it gives the largest lead distance of all the EWS), red tiles represent a minimum (that is, the smallest lead distance of all EWS), yellow tiles indicate that all EWS have the same value, grey tiles represent neutral lead distances (neither maximum nor minimum). Black tiles represent failed warnings (negative lead distances), and white tiles represent indeterminate values. Infectivity is $p = 0.2$ for panels on the left, and $p = 0.8$ for panels on the right.

Model	EWS	$\langle \mathcal{M} \rangle$	$\langle \mathcal{C} \rangle$	$\langle \mathcal{I} \rangle$	$\langle N, N \rangle$	$\langle N, V_s \rangle$	$\langle V_s, V_s \rangle$
V2	SNHT	+41	-13	-13	-22	+27	+27
	Lan	+36	-13	-27	-27	+27	+18
	Pet	+46	-22	-4	-64	+55	+45
	BR	+36	+4	+18	-32	+14	+4
V1 $p = 0.2$	SNHT	+36	-16	-12	-55	+28	+26
	Pet	+38	-20	-10	-61	+56	+54
	Lan	+35	-15	-6	-49	+36	+45
	BR	+32	-10	-19	-61	+37	+25
V1 $p = 0.8$	SNHT	+48	-14	-15	-65	+52	+41
	Pet	+40	-26	-15	-45	+52	+66
	Lan	+40	-26	-15	-65	+67	+78
	BR	+41	-28	-18	-66	+48	+48

Table A.2: Table showing the differences of the proportions of maxima and minima for each EWS, with respect to model and change point test. Green-highlighted cells represent the change point detection test giving the best performance of each EWS. Choices of best EWS were not made when all results were negative.

A.2.8 A grand comparison

Taking only the general performance of the EWS into account, a metric can be derived from Figs. A.19 and A.20 by simply subtracting the proportion of minima from the proportion of maxima for each EWS, per change point test and model. The resulting differences can be seen in Tab A.2. Green-highlighted cells serve to indicate the change point detection test corresponding to the best performance of the EWS. Where an EWS has equal scores under different tests, preference is given to the test that resulted in less minima (instances of smallest lead distance). Note that ‘best tests’ were not chosen for for the anti-vaccine similar join count $\langle N, N \rangle$ in all models because all scores were negative; this signals fallibility of this EWS, and therefore we discourage its use. This is also the reason why a best test was not chosen for Geary’s C $\langle \mathcal{C} \rangle$ and Moran’s I $\langle \mathcal{I} \rangle$ for model V1. The poor performance of the anti-vaccine similar join count continues the theme of the Results (2.3), whereas the confinement of the validity of Geary’s C and Moran’s I to the larger model V1 gives us pause as related to future use.

A.3 Summation

In sum, the analysis and comparison of the results of these models allow the conclusion that the behaviours of the EWS are largely independent of the parameters that differed between the models. Our main results were the reconfirmation of $\langle N, V_s \rangle$ and *mutual* as the leading EWS of the study through their high lead distances, a new attention paid to the use of $\langle N, N \rangle$ as an indicator of transition in V1, and the confirmation of vanishing intertransition distance inherent to the model (a cause for concern, as stated in Sec. 2.4 (Discussion)). However, it appears that individual lead distance of each EWS is subject to the size of the network, the change point detection tests used (each pros and cons [175]), and the (method of) application of these tests for determining the exact location of the first warning given by each EWS.

Appendix B

Community structure predicts social shifts

B.1 Lead distance plots

The figures in this section show the lead distances given by various EWS with respect to the social norm σ with different disease infectivities p and the four different change point tests. In all figures: panels (a) represent the standard normal homogeneity test (SNHT), panels (b) represent the Lanzante test, panels (c) the Pettitt test and panels (d) the Buishand range test. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

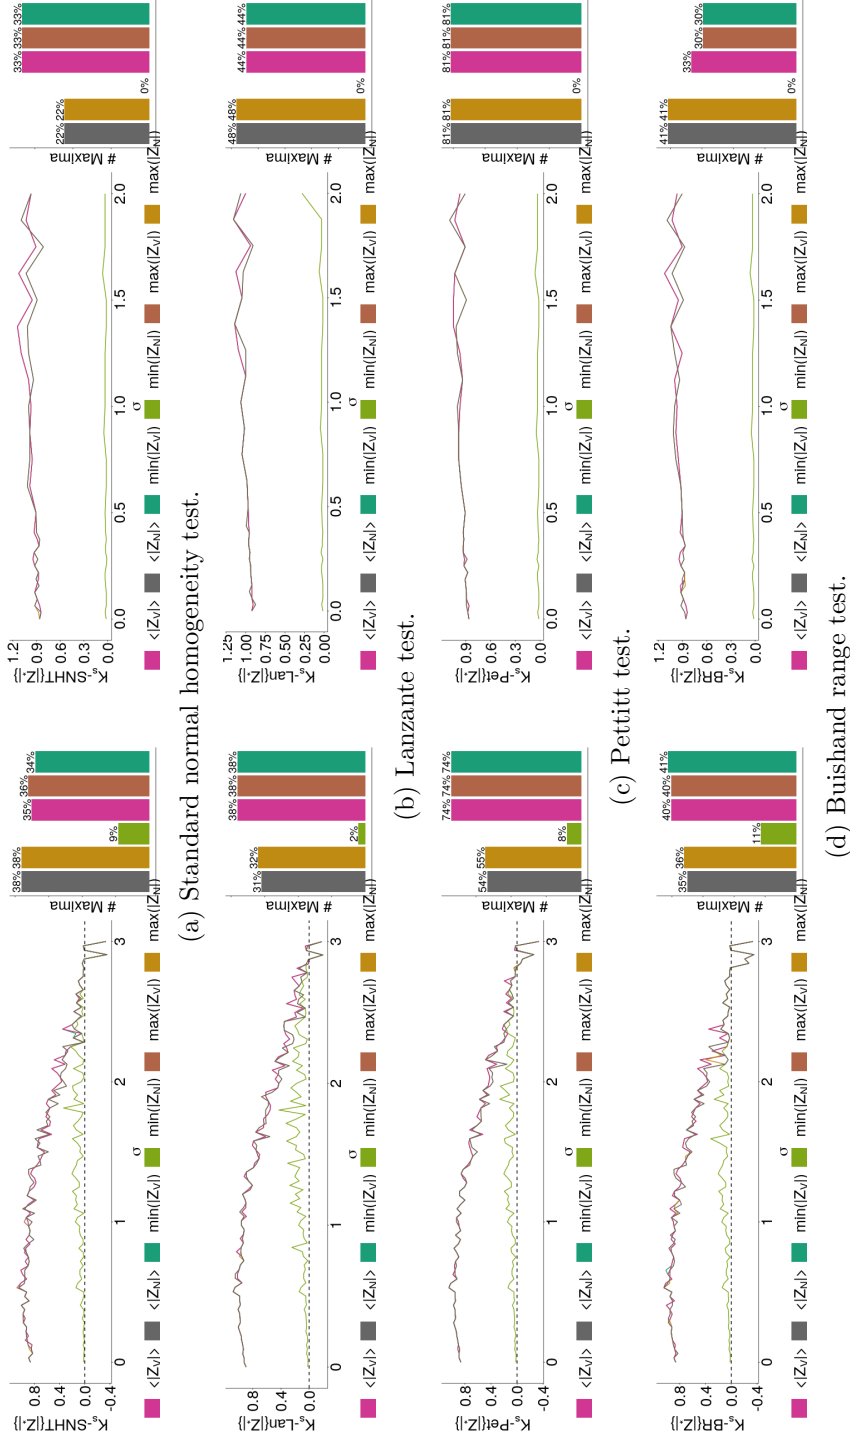


Figure B.1: Trends in the lead distances of the sizes of different types of opinion communities ($|Z_{*}|$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

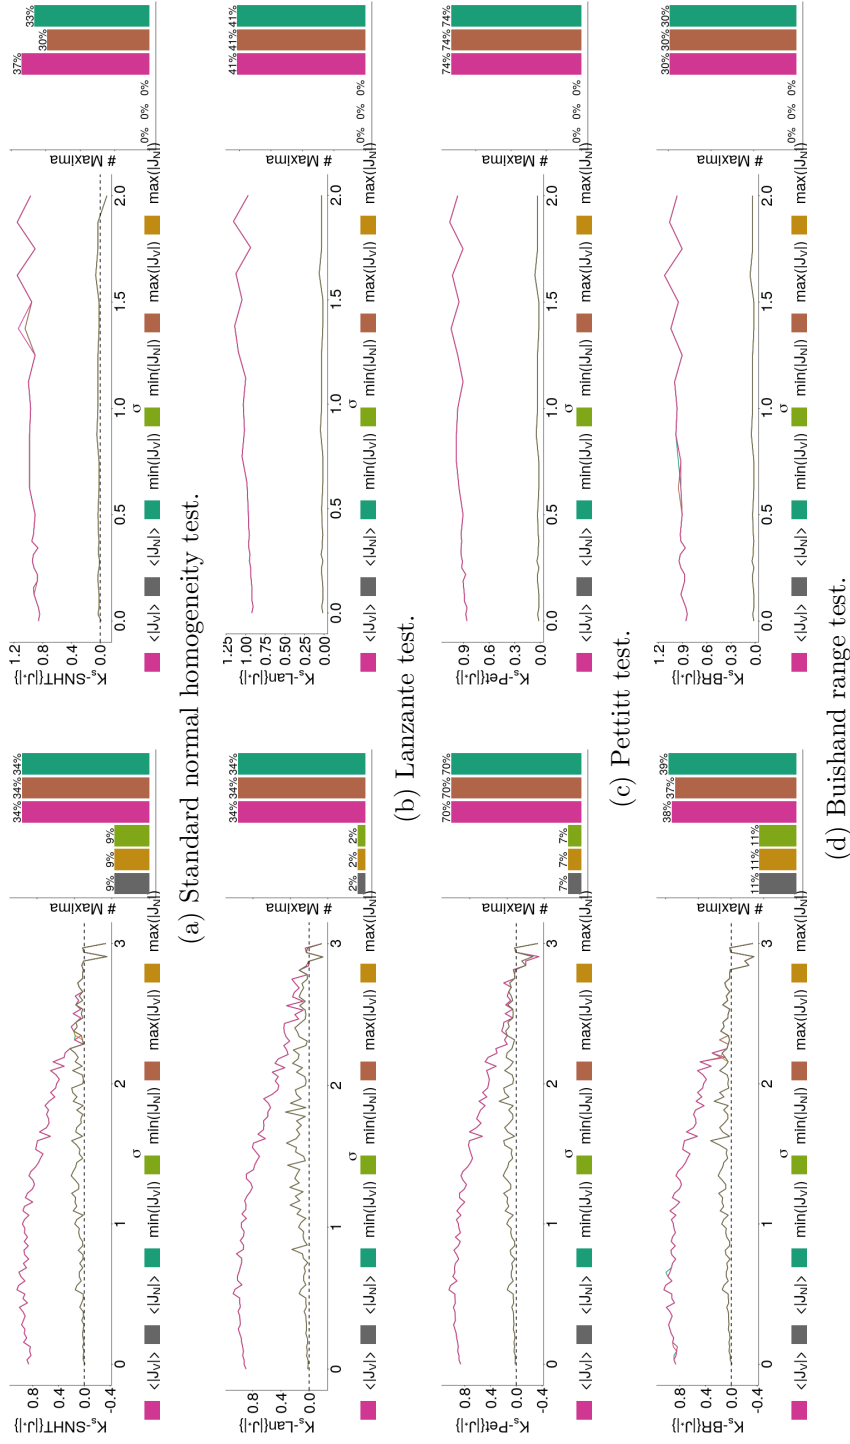


Figure B.2: Trends in the lead distances of different types of echo chambers ($|J^*|$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

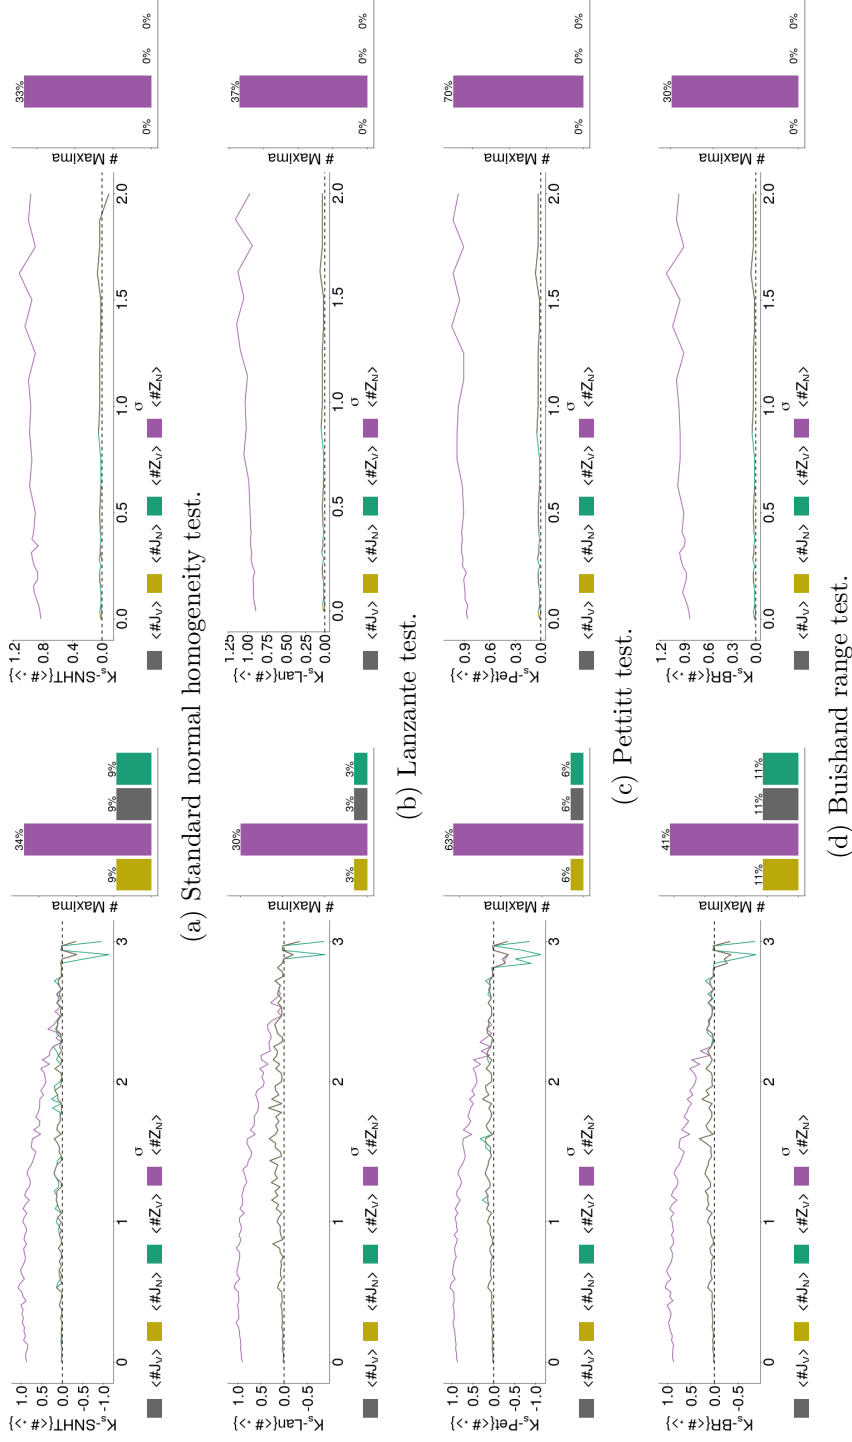


Figure B.3: Trends in the lead distances of the censuses of different opinion communities ($\#Z_*$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

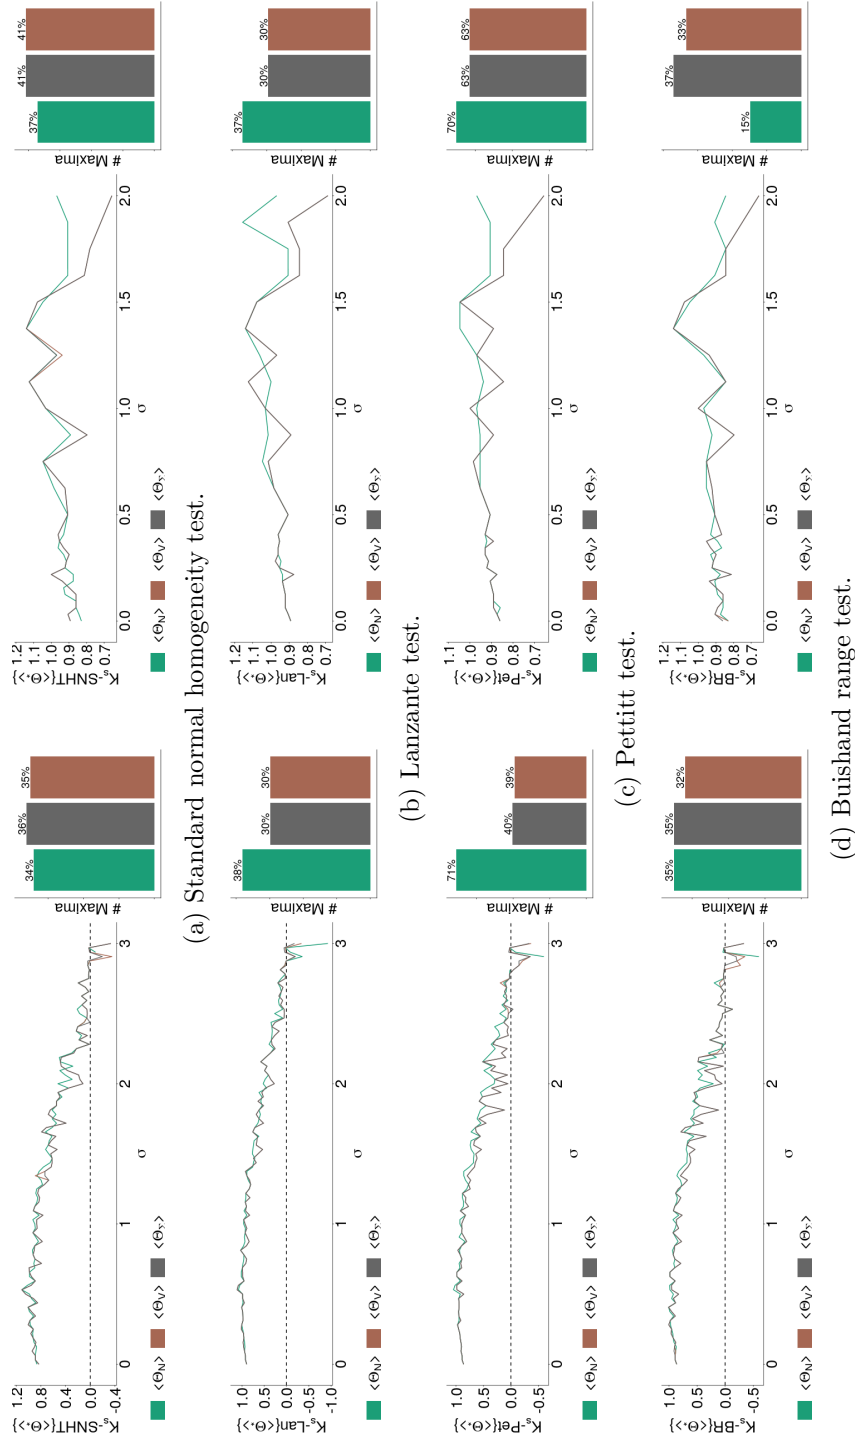


Figure B-4: Trends in the lead distances of the number of different opinion changes ($\langle \Theta_* \rangle$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

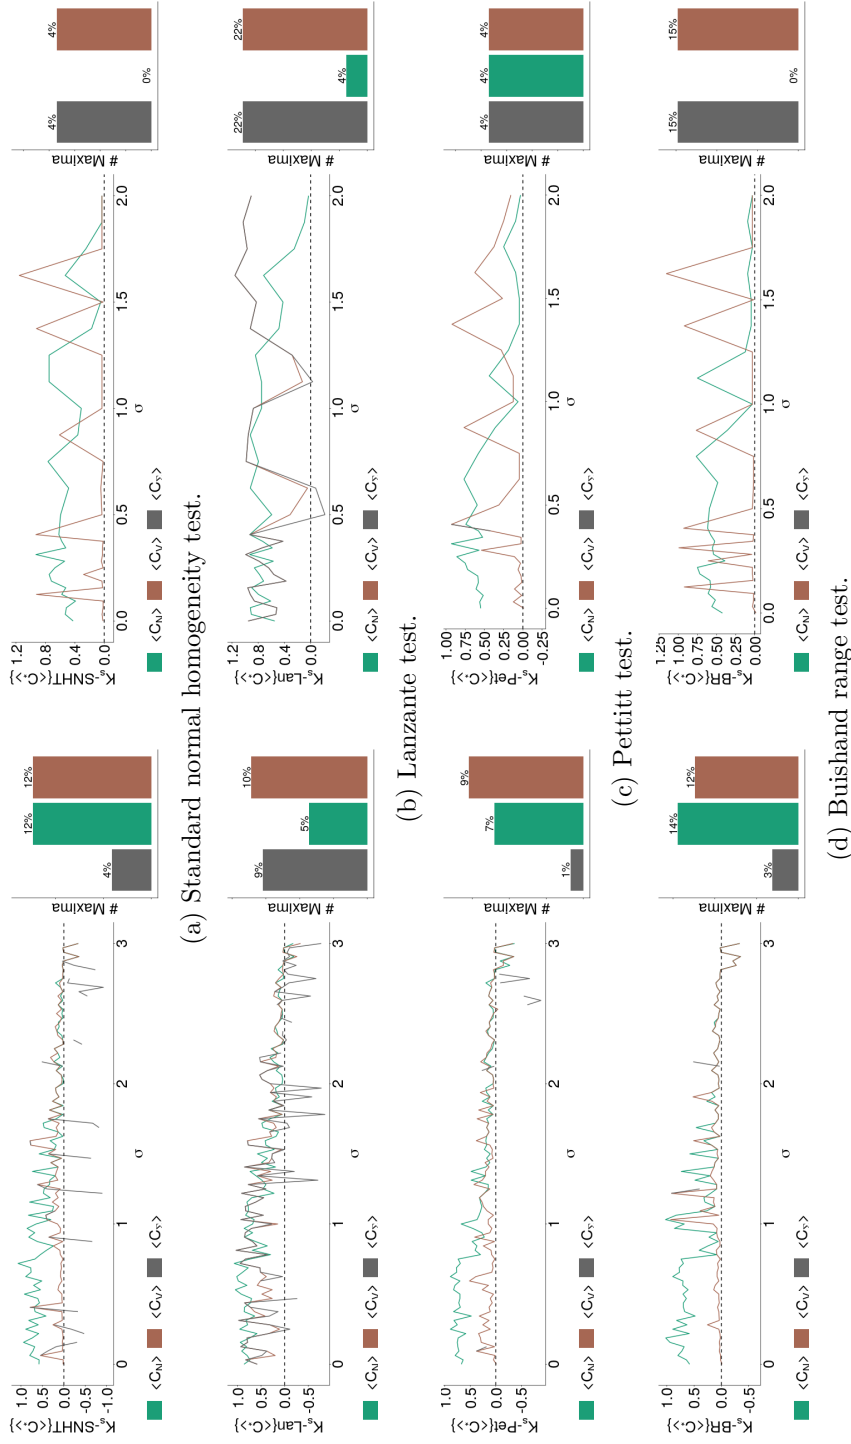


Figure B.5: Trends in the lead distances of the global clustering coefficient ($\langle C_* \rangle$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

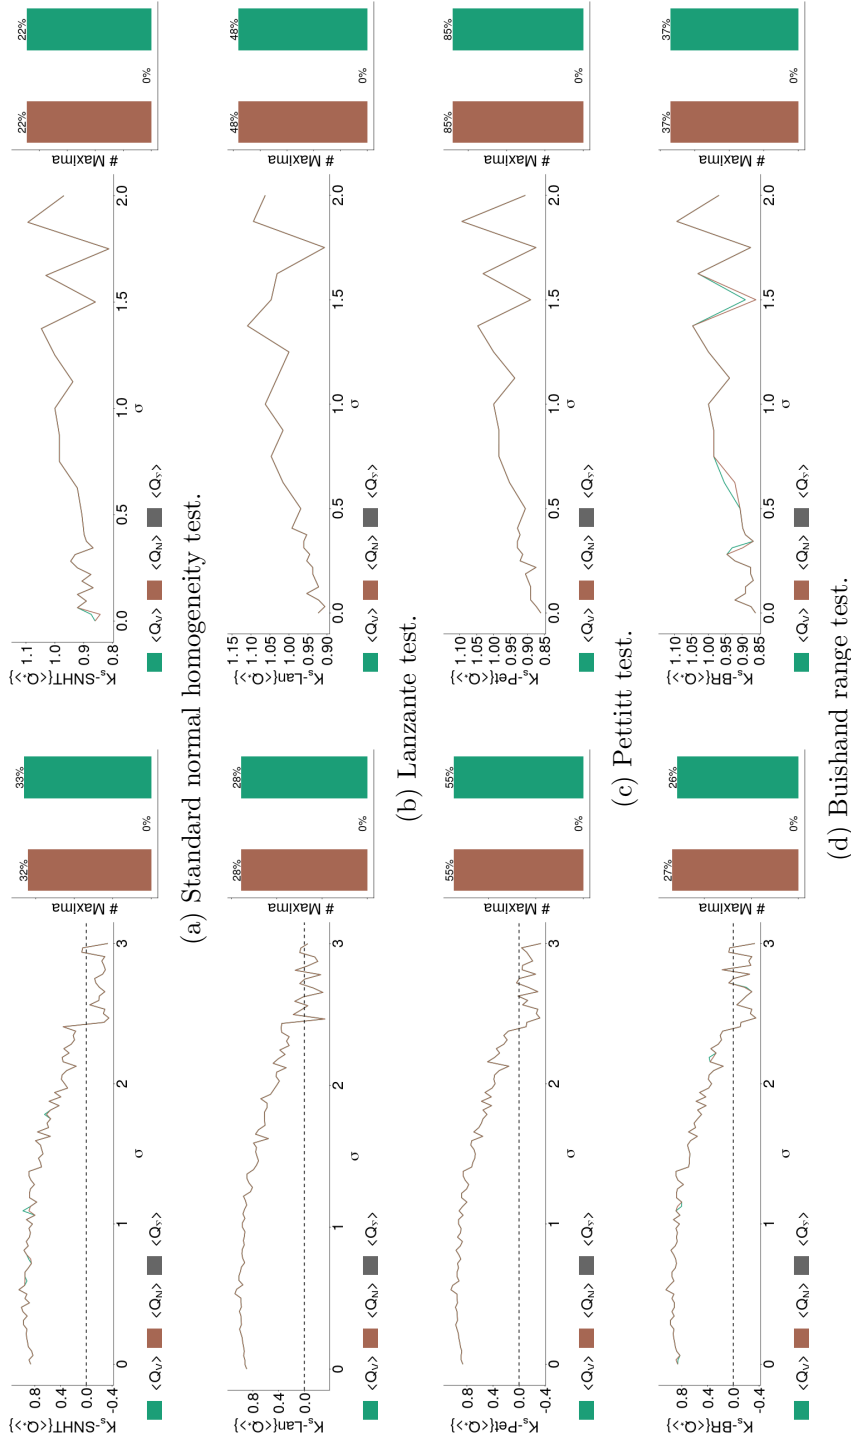


Figure B.6: Trends in the lead distances of the modularity statistic (Q_*) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

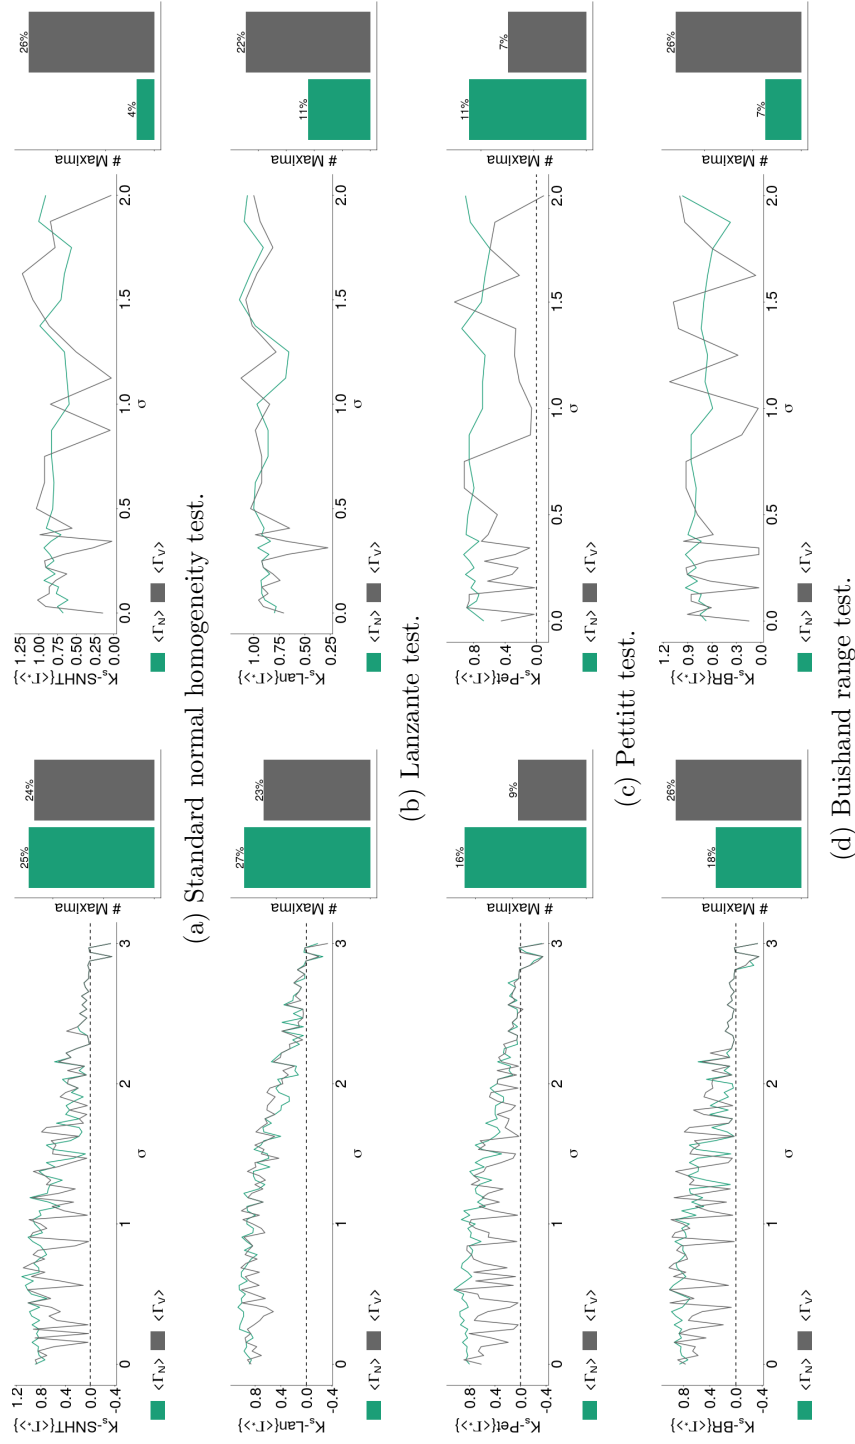


Figure B.7: Trends in the lead distance given by the probability of having an infected neighbour ($\langle \Gamma_* \rangle$) under different change point tests. Left panels give results for disease infectivity $p = 0.2$, while panels on the right represent infectivity $p = 0.8$.

B.2 EWS performance grid plots

All the figures in this section are grids showing the relative performance of each EWS with various change point detection tests and disease infectivities. Green tiles denote the social norm (σ) values for which the EWS gave the highest lead distance, red tiles represent the smallest lead distance, grey tiles represent lead distances that are neither maxima nor minima and yellow tiles show where all EWS gave the same lead distances. Black squares represent failed warnings (negative lead distances) and white tiles represent undefined values (no lead distance). In all figures in this section, panels (a) represents the infectivity $p = 0.2$ and panels (b) represent the infectivity $p = 0.8$.







Figure B.8: Performance of each EWS per σ value with the SNHT.




(a) Infectivity $p = 0.2$.




(b) Infectivity $p = 0.8$.

Figure B.9: Performance of each EWS per σ value with the Lanzante test.




(a) Infectivity $p = 0.2$.




(b) Infectivity $p = 0.8$.

Figure B.10: Performance of each EWS per σ value with the Pettitt test.



(a) Infectivity $p = 0.2$.



(b) Infectivity $p = 0.8$.

Figure B.11: Performance of each EWS per σ value with the Buishand range test.

B.3 EWS performance bar charts

Grid showing the absolute and relative performance of each EWS. Green tiles denote the social norm σ values for which the EWS gave the highest lead distance, red tiles represent the smallest lead distance, grey tiles represent lead distances that are neither maxima nor minima and yellow tiles show where all EWS gave the same lead distances. Black squares represent failed warnings (negative lead distances) and white tiles represent undefined values (no lead distance). In all figures in this section, panels (a) represents the infectivity $p = 0.2$ and panels (b) represent the infectivity $p = 0.8$.




Figure B.12: Grids showing the absolute and relative performance of the EWS for different infectivities under the SNHT.




Figure B.13: Grids showing the absolute and relative performance of the EWS for different infectivities under the Pettitt test.




Figure B.14: Grids showing the absolute and relative performance of the EWS for different infectivities under the Lanzante test.




Figure B.15: Grids showing the absolute and relative performance of the EWS for different infectivities under the Buishand range test.

Appendix C

Spatial EWS are robust to delay and network sampling

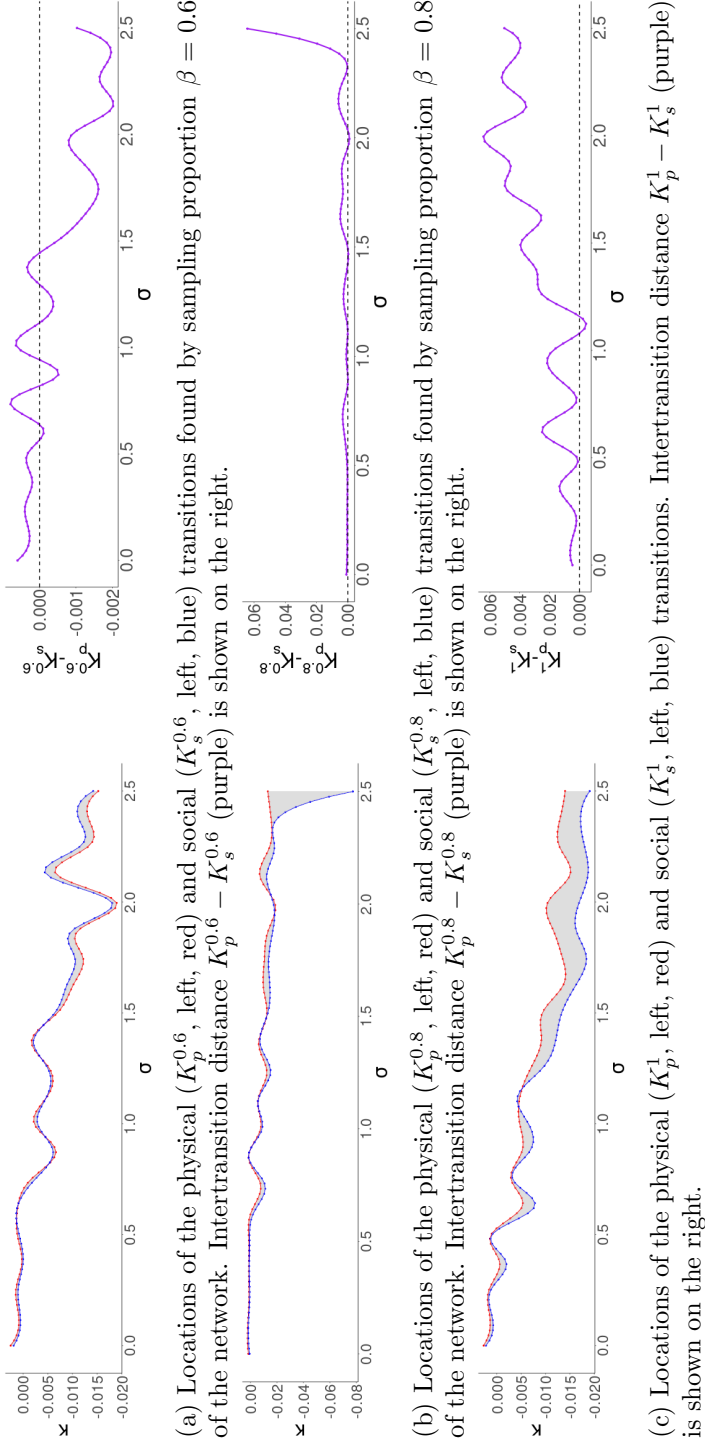


Figure C.1: Locations of the physical and social transitions (K_p^β and K_s^β), and trends of the intertransition distance $K_p^\beta - K_s^\beta$ obtained by sampling different proportions β of agents in the network.

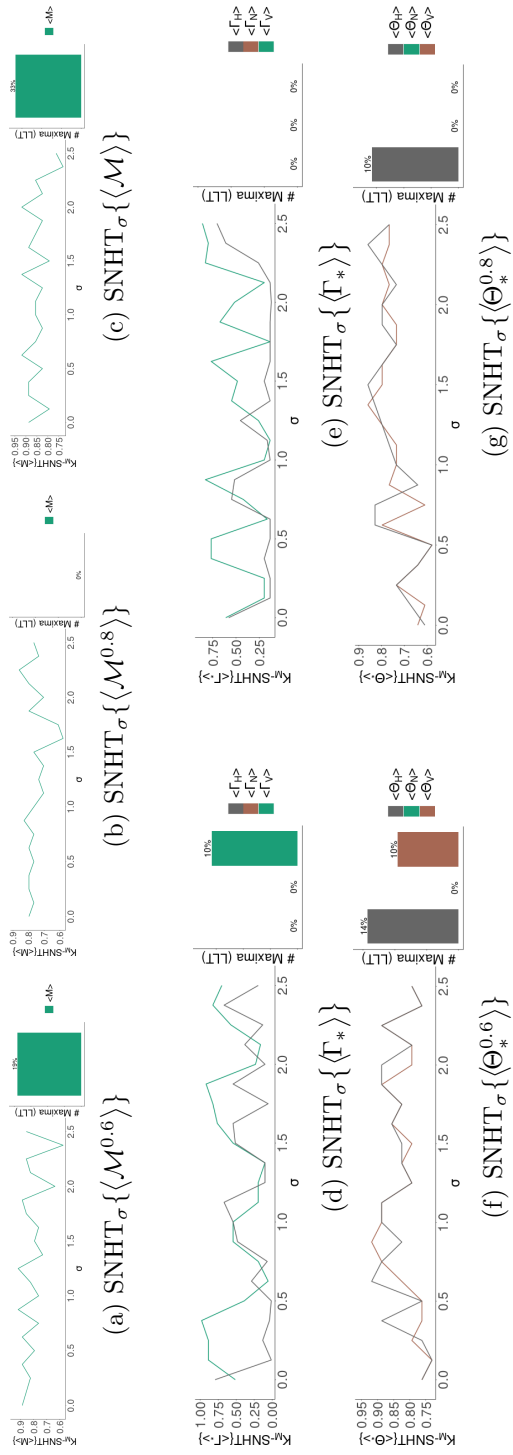


Figure C.2: Changes in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (1).

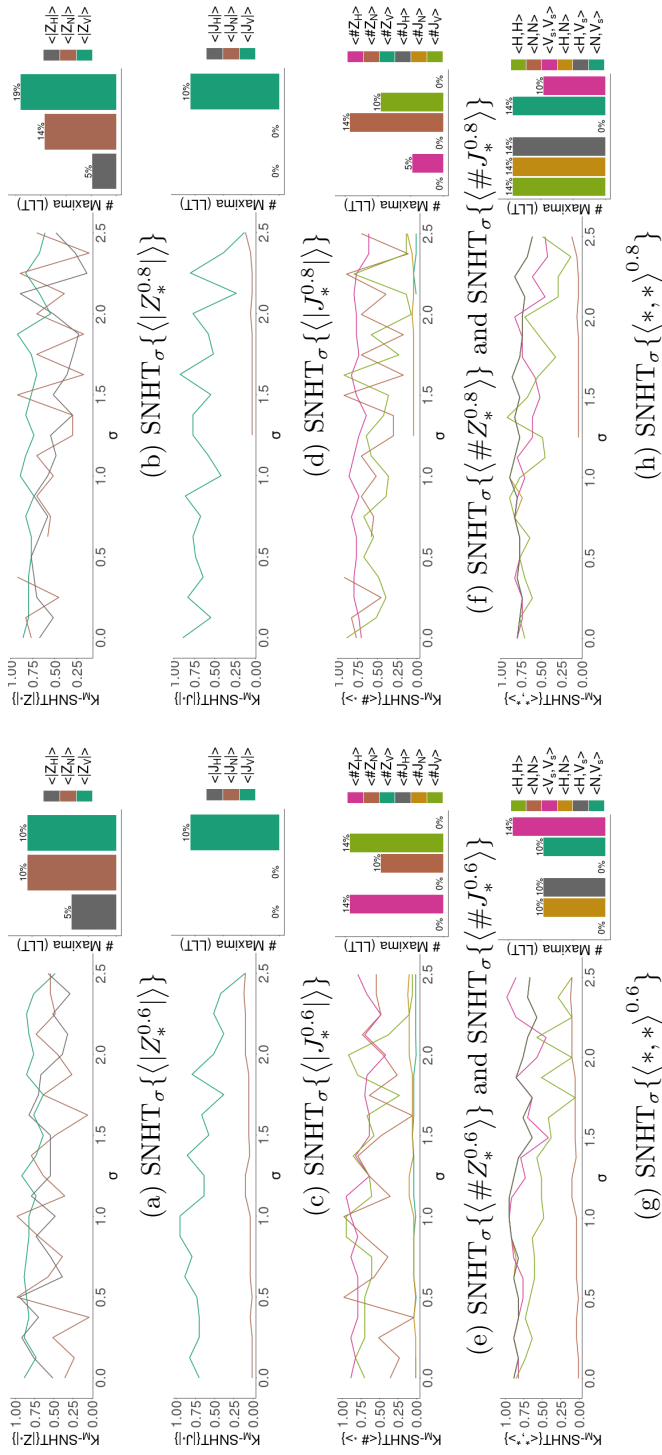


Figure C.3: Changes in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (2).

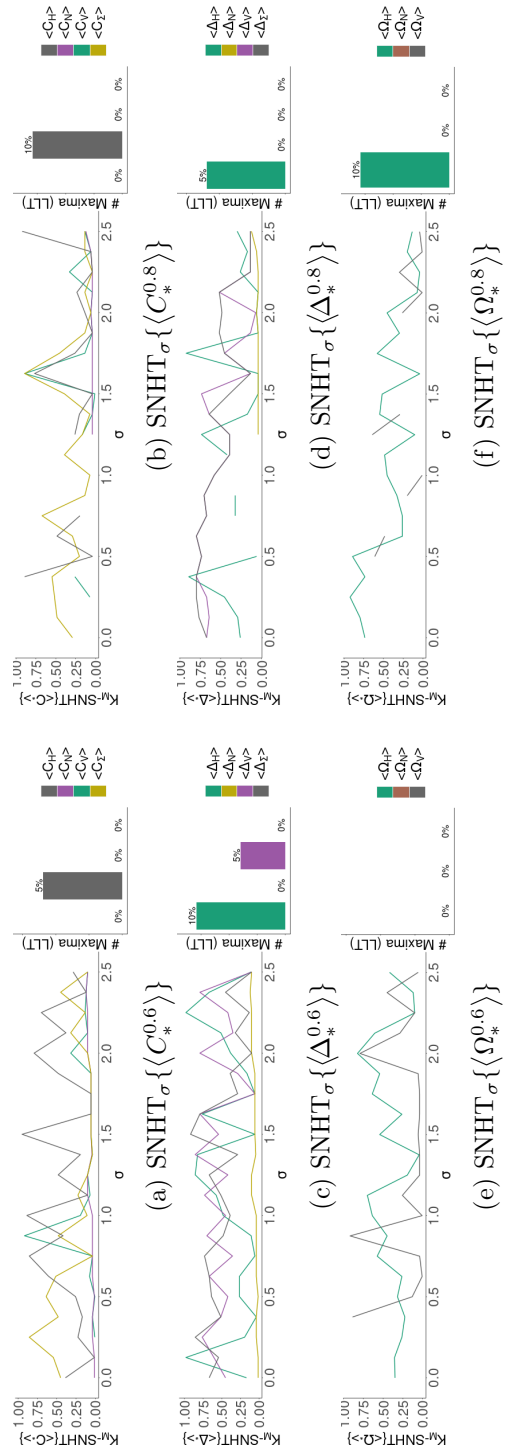




Figure C.4: Variance in the trends of the lead times of some EWS caused by sampling a proportion β of the social network in model V4 (3).



(a) Grid showing the performance of EWS for each value of the social norm σ .



(b) Bar chart showing the summary performance of EWS.

Figure C.5: Performance of the EWS with respect to the value of the social norm σ when 60% of the network is sampled.

Appendix D

Model-based projections for COVID-19 outbreak size

D.1 Sensitivity Analysis

D.1.1 Varying α_0 and B_H

The parameter β^H represents the rate of interaction in the household, and thereby regulates the spread of the disease. For each value of α_0 , increasing the rate of interaction in the home β^H increases the number of infections produced for both RA (Supplementary Fig. D.1) and ST (Supplementary Fig. D.2) allocation. In most scenarios (7:3 RA being one of the exceptions), varying α_0 (for constant β^H) produces a small increase in the number of infections produced throughout the simulation. The rate of increase also depends on the number of children in the classroom; for the scenario 31:1 RA, increasing β^H from 0.0545 to its baseline value 0.109 almost triples the number of total infections.

D.1.2 Varying α_0 and R_{init}

The parameter R_{init} refers to the proportion of individuals we presume are recovered from some previous period of infection spread, while α_0 is responsible for the rate of infection in common areas relative to the infection rate in the classroom. All other parameters are set to the baseline values given in Supplementary Tab. 5.2. These parameters were varied together by 50% in either direction. In Supplementary Figs. D.3 and D.4, increasing values of R_{init} lower both the means and standard deviations of the total number of infections for each value of α_0 . Also, for each value of R_{init} , the total number of infections produced increases with α_0 . This shows opposing




Figure D.1: Results of varying the parameters β^H and α_0 by (50% each) on the total number of produced infections for RA allocation. Error bars denote a single standard deviation of the data used, and boxed text shows the corresponding mean and standard deviation.




Figure D.2: Results of varying the parameters β^H and α_0 by (50% each) on the total number of produced infections for ST allocation. Error bars denote a single standard deviation of the data used, and boxed text shows the corresponding mean and standard deviation.

interaction between increasing common area infection and increasing initial recovery rate; one increases infection and the other lowers it (respectively).

D.1.3 Varying α_0 and λ_i

From Tab. 5.2, parameter λ_i varies the amount of community infection in the model (infection due to other sources not modelled, such as public transport); be reminded that we assumed that the rate of community infection is effectively twice the baseline value for those individuals in the model not attending the school.

For each value of α_0 in Supplementary Fig. D.6, the total number of infections produced in the simulation increases with λ in each scenario with random allocation (RA), and also with grouping by household (ST, Supplementary Fig. D.5). For each λ , there is no consistent relationship between the numbers of infections and the value of α_0 . This result is intuitive; though the effect is not pronounced, increasing the rate of community infection increases the total number of infections in each tested scenario.




Figure D.3: Results of varying the parameters R_{init} and α_0 by (50% each) on the total number of infections for ST allocation. Text in boxes denotes the mean and standard deviation of the data corresponding to the parameters and error bars denote a single standard deviation of the data used.




Figure D.4: Results of varying the parameters R_{init} and α_0 by (50% each) on the total number of infections for RA allocation. Text in boxes denotes the mean and standard deviation of the data corresponding to the parameters and error bars denote a single standard deviation of the data used.




Figure D.5: Results of varying the parameters λ_i and α_0 by (50% each) on the total number of infections for ST allocation. Text in boxes denotes the mean and standard deviation of the data corresponding to the parameters and error bars denote a single standard deviation of the data used.




Figure D.6: Results of varying the parameters λ_i and α_0 by (50% each) on the total number of infections for ST allocation. Text in boxes denotes the mean and standard deviation of the data corresponding to the parameters and error bars denote a single standard deviation of the data used.

Glossary

Agent-based model (ABM) A computational model of a process where discrete units with individual attributes (called *agents*) are allowed to interact with each other and some encompassing environment within certain imposed requirements. Unlike EBMs, ABMs are perceived as being more ‘realistic’ due to the elimination of the *well-mixing* and *homogeneity* assumptions. However, due to their complexity, these models usually don’t allow closed-form solutions, instead yielding approximations of key statistics.

Assortative mixing A pattern of connection on a [network](#) in which agents choose to connect to other agents with a shared characteristic (in this case, identical vaccine opinion).

Basic reproductive ratio Denoted \mathcal{R}_0 , this represents the expected number of infections directly resulting from initial infected case in a population of susceptible individuals in the absence of public health intervention. Also called the **basic reproductive number**.

Bistability The existence of two stable equilibria in a dynamical system.

Case importation The process by which infectious cases arrive in the studied region from elsewhere.

Change point A member of a set of definable points in a time series, stochastic process or other ordered sequence of values at which the probability distribution changes significantly.

Change point detection test These are statistical tests used to find [change points](#) in time series or ordered series of values. In this thesis, we use the Lanzante, Pettitt, Buishand range tests (Chs. [2](#) and [3](#)) and [SNHT](#) (Ch. [4](#)). We also restrict investigation to the first change point found in a κ -series.

Cohort A group of students that will have face-to-face teaching and online classes in alternating weeks.

Community An interconnected group of like-minded agents, here treated as a connected component of agreeing agents in the network.

- Community spread** The spread of infection in a region without known contact with some infected person.
- Contact tracing** The process of identifying and tracking individuals who have been exposed to an infected case, usually with the aim of slowing or preventing further infection and disease transmission.
- Core** Describes an agent in an opinion [community](#) whose neighbours all share their opinion. A group of core agents is mainly referred to as an [echo chamber](#).
- Correlation length** The length over which two microscopic variables are related on a thermodynamic system. Specifically, let x and y be two lattice sites and $s(x)$ and $s(y)$ be their spins (respectively). Correlation length is given by the term $\langle s(x) \cdot s(y) \rangle$, giving the ‘strength’ of the correlation between the spins at sites x and y .
- COVID-19** A disease caused by the SARS-CoV-2 coronavirus. Commonly called simply *coronavirus*, it was first identified in December 2019 in China and has since spread to all countries around the world.
- Critical point** The point at which a physical system undergoes a [second-order phase transition](#).
- Critical slowing down** The phenomenon by which the relaxation time (the time taken by the system to return to equilibrium after some small perturbation) increases as a [critical point](#) is approached.
- Critical transition** A (sudden) transition when a physical system undergoes a [second-order phase transition](#) around a [critical point](#).
- Early warning signal (EWS)** A statistically significant, recognisable and characteristic behaviour known to precede [critical transition](#) in dynamical systems.
- Echo chamber** Generally defined as a well connected groups of people promoting and reinforcing the same bias, or alternately, a closed subnetwork within a larger structure. In Ch. 3, it refers to a set of mutually connected agents with only like-minded neighbours. Also called a **tribe**. Demonstrated in Fig. 3.1.
- Effective contact** An interaction between agents that allows for interaction and/or communication of a vaccine opinion.
- Effective interaction** See [effective contact](#).
- Effective reproduction number** The actual number of secondary infections produced by some primary case. Represented by \mathcal{R}_e .

Epidemic A loosely defined term referring to a rapidly spreading infection. Discussed at length and defined in Sec. 1.2.

Equation-based model (EBM) A system of deterministic (ordinary and/or partial) differential equations used to model the spread of disease in some population. The simplicity of the models sometimes allows for the derivation of exact symbolic results such as the basic reproductive ratio \mathcal{R}_0 . However, such models inherently carry heavy assumptions such as *homogeneity* (individuals are assumed to be identical).

Equilibrium In this thesis, we state that a simulation has reached (*computational*) *equilibrium* when the model's output variables over the last 500 time steps have a standard deviation of less than 0.05% of their respective maximum values. Stated in Sec. 1.4.

Erdős-Rényi An algorithm for the construction of random graphs $\mathcal{G}(N, p)$ where an edge is connected between each distinct pair of nodes in the graph with probability p . The value of p controls the number of connected components of the graph; the formation of the graph is similar to a bond percolation of the complete graph K_N .

Filter bubble The isolation and reinforcement of ideas stemming from the personalisation of results given by algorithms and online search engines.

Geary's C A measure of spatial correlation using the difference in opinion score between connected agents.

Global clustering coefficient (GCC) A graph-theoretic measure of clustering describing the number of closed triplets in the network.

Graph diameter The longest distance between two vertices present in a graph, or equivalently the maximum eccentricity of the vertices of all the nodes in a graph.

Hub A massively popular node, disproportionately likely to be connected to a large number of nodes, usually not reciprocally.

Intertransition distance The gap $K_p - K_s$, where K_s and K_p represent transition points in the social and infection dynamics respectively.

Join count The number of connections between patches (locations) of different types in a spatial system. Usually the patches are placed into two attributive classes, generally termed black (B) and white (W). The number of black-white joins will be denoted $[B, W]$, for example.

Lead distance The gap $K_s - \Pi_\Psi$, where Π_Ψ represents some κ value at which we can assert that a signal occurs in some sequence Ψ of κ values (usually a [change point](#) obtained from some [change point detection test](#)). Defined in Ch. 2 as $K_s - \text{Lan}_\sigma\{\Psi\}$ and in Chs. 3 and 4 as $K_s - \text{SNHT}_\sigma\{\langle\Psi\rangle\}$.

Maximax The strategy of choosing the ‘best of the best’ option. In this case, a maximax decision would be $\max(\max(\lambda))$, where $\lambda = K_* - \Pi_\Psi$ is some [lead distance](#).

Maximin The strategy of choosing the ‘best of the worst’ option. In this case, a maximin decision would be $\max(\min(\lambda))$, where $\lambda = K_* - \Pi_\Psi$ is some [lead distance](#).

Meme A piece of information that is spread from person to person within a culture [146]. On social media, this term usually refers to an image, usually with superimposed, non-informative text.

Model variables A model variable is any of the outputs $\langle S \rangle$ (mean number of susceptible agents), $\langle I \rangle$ (mean number of infected/infectious agents), $\langle R \rangle$ (recovered agents), $\langle V_p \rangle$ (vaccinated agents), $\langle N \rangle$ (anti-vaccine agents), $\langle H \rangle$ (vaccine hesitant agents) and $\langle V_s \rangle$ (pro-vaccine agents).

Modularity A graph-theoretic quantification of the degree of segmentation present in a [network](#). Denoted Q .

Moran’s I A quantification of spatial correlation measuring the correlation between values of adjacent sites/locations in a spatial structure. Denoted \mathcal{I} .

Mutual Information An entropy-based quantification of the ‘*shared information*’ between two processes. Specifically, it measures the increase in certainty gained about one process by monitoring another. Denoted \mathcal{M} .

Network Alternately a system of interacting individuals or the structure of such interconnectivity. Usually represented as a [graph](#), where the agents in the population are represented as nodes and edges represent any effective contact between agents in the model dynamics.

Performance penalty A quantification [Pen](#) of the effect of network sampling on the performance of [EWS](#) with respect to the proportion β of the network sampled.

Peripheral Describes an agent in an opinion [community](#) that is connected to at least one neighbour with a differing opinion.

Phase transition The process by which a physical system moves from one stable state to another. Phase transitions can either be discontinuous (*first-order*) or continuous (*second-order*) depending on the dynamics of the model studied.

Preferential attachment A process by which some attribute is distributed among a set of individuals according to the amount already accrued individually. In networks, this is the process by which new agents added to the network are more likely to become connected to other agents with already large neighbourhood sizes.

Random assignment (RA) The classroom assortment strategy where classes are filled randomly (with no concern for the students' family groups).

Realisation A single instance of a simulation with a specific set of parameter values, with the collection of all realisations called an **ensemble**. Most results presented in this thesis are *ensemble means*, representing the average of the results given by all realisations run with the desired set of parameter values. Our ensemble sizes range from 20 to 50 (so that most data points presented here are the means of at least 20 values).

Scale-free network A theoretical model of a preferentially attached network whose degree distribution (asymptotically) follows a power law.

Second-order Second-order [phase transitions](#) occur when some microscopic variable of the system varies continuously (with a discontinuous second derivative), unlike *first-order transitions* where a microscopic variable varies *discontinuously* (discontinuous first derivative).

Secondary attack rate The probability of an infected case occurring within some specific group (such as a family unit).

Siblings together (ST) The classroom assortment strategy where student are assigned to classrooms in a way that keeps siblings together (as far as possible).

Small world A theoretical network structure characterised by high degrees of local clustering and small diameter. More rigorously, the expected network distance between two nodes L varies with the logarithm of the network size D , so that $L \propto \ln N$. Also referred to as a **small world network**.

Triad A pairwise connected group of three (3) agents (triplet, K_3) with the same opinion.

Triad census A count of the connected opinion [triads](#) of agents in the (social) [network](#).

Triadic closure The process of creating a triad from a dyad by connecting the two unconnected members. Demonstrated in Fig. 4.8.

Vaccine "A suspension of live (usually attenuated) or inactivated microorganisms (e.g. bacteria or viruses) or fractions thereof administered to induce immunity and prevent infectious diseases and their sequelae." [101]

Vaccine hesitancy “... delay in acceptance or refusal of vaccination despite availability of vaccination services.” [410]

Watts-Strogatz An algorithm for the construction of random graphs $\mathcal{WS}(N, K, \beta)$, where a ring lattice is first constructed by joining each node to its $\frac{K}{2}$ left and right neighbours. Each edge connecting each node to its right neighbours is then rewired with probability β without edge duplication. The produced graphs feature **small world** properties such as a high degree of clustering and small average path lengths. $\beta = 0$ gives a ring lattice, while $\beta = 1$ gives an **Erdős-Rényi** graph.