

BIO-SURVEILLANCE: DETECTION AND MITIGATION OF DISEASE OUTBREAKS

A Thesis
Presented to
The Academic Faculty

by

Mi Lim Lee

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Industrial and Systems Engineering

Georgia Institute of Technology
December 2013

Copyright © 2013 by Mi Lim Lee

BIO-SURVEILLANCE: DETECTION AND MITIGATION OF DISEASE OUTBREAKS

Approved by:

Professor David Goldsman, Advisor
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor Seong-Hee Kim, Advisor
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor Kwok-Leung Tsui, Advisor
Department of Industrial Engineering
City University of Hong Kong

Professor Sigrun Andradottir
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor Roshan Joseph Vengazhiyil
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Date Approved: 6 September 2013

To my awesome husband and family.

ACKNOWLEDGEMENTS

Pursuing my Ph.D. degree would be remembered as the most challenging but exciting journey in my life. Every moment of the journey has been shared with many great people and I would like to thank them for their support.

I would like to express the deepest appreciation to my advisors, Dr. David M. Goldsman, Dr. Seong-Hee Kim, and Dr. Kwok-Leung Tsui. I have learned a lot from their vision, enthusiasm, and immense knowledge during my doctoral program. This dissertation would not have been finished without their guidance and encouragement.

Besides my advisors, I would like to thank the rest of my thesis committee, Dr. Sigrún Andradóttir and Dr. Roshan Vengazhiyil, for their insightful questions and helpful suggestions.

My thanks also go to my friends: Dong Gu Choi, Seonghye Jeon, Judy Sunyoung Lee, Haewon Park, Hin Kyeol Woo, THEM, and many other fellow students in the Industrial and Systems Engineering department. They were the sources of all the laughter, joy, and fun in my graduate life.

My parents, Gweon Heui Lee and Kyung Ae Kim, deserve my sincere thanks for giving me a birth at the first place. Their endless love has supported me even in the sleepless nights I was struggling with my research.

Last but not least, special thanks to my husband, Chuljin Park, and my first child, Dodo, who enrich my life. I look forward to starting a whole new journey with them after my Ph.D.

Contents

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vii
LIST OF FIGURES	ix
SUMMARY	xi
I INTRODUCTION	1
II REACTIVE STRATEGIES FOR DEVELOPING OUTBREAKS OF PANDEMIC INFLUENZA	5
2.1 Simulation Model	6
2.1.1 Disease Transmission Model	9
2.1.2 Intervention Strategies	14
2.2 Economic Cost Estimation	18
2.3 Simulation Results	22
2.4 Discussion	29
III SPATIOTEMPORAL BIOSURVEILLANCE: CONTROL LIMIT APPROXIMATION AND THE IMPACT OF SPATIAL CORRE- LATION	32
3.1 Background	36
3.2 Design of the JMCUSUM Chart	39
3.2.1 Fixed Scan Radius	40
3.2.2 Variable Scan Radius	41
3.3 Approximation Accuracy	43
3.3.1 Experimental Setup	43
3.3.2 Range of the Scan Radius	44
3.3.3 Results for the Fixed Scan Radius Case	48
3.3.4 Results for the Variable Scan Radius Case	50

3.4	Impact of Spatial Correlation on Detection Performance	51
3.4.1	Outbreak Radius	54
3.4.2	Spatial Correlation	54
3.5	Conclusions	55
IV	ROBUST DISTRIBUTION-FREE MCUSUM CHARTS FOR SPA- TIOTEMPORAL BIOSURVEILLANCE IN THE PRESENCE OF SPATIAL CORRELATION	57
4.1	Background and Motivation	58
4.1.1	Notation and Problem	58
4.1.2	Motivation	61
4.2	RMCUSUM Charts	63
4.2.1	RMCUSUM Chart for the Fixed Scan Radius	64
4.2.2	RMCUSUM Chart for the Variable Scan Radius	66
4.3	Experiments	72
4.3.1	Experimental Setup	72
4.3.2	Accuracy of Control Limits	74
4.3.3	ARL ₁ Performance	77
4.4	Conclusions	79
V	CONTRIBUTIONS	80
	APPENDIX A	82
	APPENDIX B	86
	APPENDIX C	90
	REFERENCES	91

List of Tables

1	Per-contact influenza transmission probabilities within contact groups	11
2	Age-group-specific illness attack rates and calibrated attack rates . . .	12
3	Attack rates and total costs of interventions with 35% vaccination coverage	16
4	Attack rates and total costs of interventions with 70% vaccination coverage	17
5	Proportions of influenza cases at high risk for complications ¹	19
6	Outpatient visit, hospitalization, and death rates ¹	19
7	Outpatient visit, hospitalization, and death costs ¹	21
8	Costs (in US\$) and impacts of vaccination ¹	22
9	r versus k when $\delta = 1$ and $\rho = 0$	45
10	Empirical h^{**} (Analytical h^{**}) when r is fixed	49
11	Empirical ARL_0 of the JMCUSUM-F chart with analytical h^{**}	49
12	Empirical ARL_1^{JF} (Analytical ARL_1^{JF}) with analytical h^{**}	50
13	Empirical ARL_0 for the SMCUSUM chart with analytical h^{*r}	51
14	Empirical ARL_1^{JV} (ARL_1^S)	51
15	Identification accuracy CP (SP) with Σ_1	53
16	Identification accuracy CP (SP) with Σ_2	53
17	Configurations	70
18	Empirical ARL_0 from the multivariate normal data	75
19	Empirical ARL_0 from the multivariate Poisson data	76
20	ARL_1 from multivariate normal data	77
21	ARL_1 from multivariate Poisson data	78
22	ARL_1 comparison	78
23	Average age-stratified overall illness attack rates (%) and cost estimates	83
24	Empirical ARL_0 of the JMCUSUM-F chart with analytical h^{**}	88
25	Empirical (analytical) ARL_1^{JF} with analytical h^{**}	89

26	Empirical ARL_0 for the SMCUSUM chart with analytical h^{*r}	89
27	Empirical ARL_1^{JV} (ARL_1^S)	89

List of Figures

1	Age distribution for simulated population	8
2	Household size distribution for simulated population	8
3	Simulation flowchart and modeled influenza natural history	13
4	Total cost of intervention strategies vs. the average illness attack rate	23
5	Daily attack rates	24
6	Illness attack rates for modeled interventions with 35% vaccine coverage	26
7	Illness attack rates for modeled interventions with 70% vaccine coverage	27
8	Total cost of modeled intervention strategies vs. average illness attack rate	28
9	Coordinate expression of regions ($p = 7 \times 7$)	37
10	Outbreak clusters when $c_{\text{out}} = 25$ (outbreak regions are in gray) . . .	43
11	Correlation settings of Σ_1 and Σ_2 for region 25	44
12	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for a single MCUSUM chart	46
13	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus k for a single MCUSUM chart	46
14	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for the JMCUSUM-F chart	47
15	Spatial correlation vs. ARL_1 and CP with $\delta = 1$ and Σ_2	55
16	Coordinate expression of regions ($p = 7 \times 7$)	59
17	The RMCUSUM_f chart	66
18	The RMCUSUM_v chart	68
19	D_ρ vs A_ρ	70
20	Estimation of A_ρ	71
21	Finding \check{u}	72
22	Outbreak clusters when $c_{\text{out}} = 25$ (outbreak regions are in gray) . . .	73
23	Correlation settings of Σ_1 and Σ_2 for region 25	74
24	Σ_3 and Σ_4	76
25	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for a single MCUSUM	87
26	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus k for a single MCUSUM	87

27	$\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for the JMCUSUM-F chart	88
----	---	----

SUMMARY

In spite of the remarkable development of modern medical treatment and technology, the threat of pandemic diseases such as anthrax, cholera, and SARS has not disappeared. As a part of emerging healthcare decision problems, many researchers have studied how to detect and contain disease outbreaks, and our research is aligned with this trend. This thesis mainly consists of two parts: epidemic simulation modeling for effective intervention strategies and spatiotemporal monitoring for outbreak detection.

We developed a stochastic epidemic simulation model of a pandemic influenza virus (H1N1) to test possible interventions within a structured population. The possible interventions — such as vaccination, antiviral treatment, household prophylaxis, school closure and social distancing — are investigated in a large number of scenarios, including delays in vaccine delivery and low and moderate efficacy of the vaccine.

Since timely and accurate detection of a disease outbreak is crucial in terms of preparation for emergencies in healthcare and biosurveillance, we suggest two spatiotemporal monitoring charts, namely, the SMCUSUM and RMCUSUM charts, to detect increases in the rate or count of disease incidents. Our research includes convenient methods to approximate the control limits of the charts. An analytical control limit approximation method for the SMCUSUM chart performs well under certain conditions on the data distribution and monitoring range. Another control limit approximation method for the RMCUSUM chart provides robust performance to various monitoring range, spatial correlation structures, and data distributions without intensive modeling of the underlying process.

Chapter I

INTRODUCTION

With the ongoing threat of pandemic diseases and bioterrorism, extensive research efforts have been undertaken on the design and performance of disease outbreak surveillance systems and mitigation strategies. Various models, originating from different research fields, have been proposed for analyzing and estimating the spread of past and future pandemics.

Epidemic simulation models that mimic events happening in the real world without making strong assumptions provide useful tools to increase our understanding of the dynamics and patterns of disease propagation. They also allow for the study and evaluation of the potential impacts of various government policies and intervention strategies for infectious diseases, including vaccination, prophylactic use of antivirals, and social distancing strategies such as school closure, quarantine, and isolation.

Over the last several years, these simulation models have been influential in the formation of pandemic preparedness plans. A simulation can explore not only the epidemiological impact but also the economic effectiveness of additional, reactive strategies for containing disease outbreaks. Some of the challenges pertaining to such simulation models can be summarized as follows.

- A simulation model must imitate the complicated real world. The population and community structure designed for the simulation should be able to capture the characteristics of a targeted population and the population's behaviors. For example, places of exposure, such as households, schools, and workplaces, should be carefully selected by observing the population and the community structures closely. To this end, a great deal of demographic input data are often required.

- The model logic and interventions depend on the builder’s purpose, social situations encountered, and disease characteristics. The builder needs to determine what kinds of interventions can be employed and how to apply them in the model. For example, the availability and amount of vaccines/antivirals can be affected by the wealth of the country and the disease dynamics. The distribution of the vaccines/antivirals can be affected by political or geographical supply chain issues. One must also study the exposure frequencies and times of individuals, such as contact rates and contact durations, as well as the disease transmission parameters such as the probability that an infected person will infect a healthy person given their contacts. For these tasks, knowledge ranging from sociology to epidemiology may be required.
- The model is carefully calibrated to predict illness attack rates to calculate costs resulting from applied interventions. This is where the validity of the simulation matters. It is difficult to track a relationship between illness attack rates and costs. In addition, both direct and indirect costs should be considered. One may find it difficult to obtain appropriate references for these costs.

A goal in this thesis is to develop a reasonable epidemic simulation model that can test possible interventions for a structured community. The model should be able to report valid estimates of illness attack rates and costs associated with the different combinations of mitigation strategies. This type of model will be quite helpful for public health authorities when reacting to future outbreaks of pandemics.

The timely and accurate detection of a disease outbreak is crucial for preparation for emergency situations. Before any studied mitigation strategies are performed, one must alarm (signal) an outbreak at an appropriate time, and decide when and where to apply the mitigation strategies.

One of the statistical approaches developed for timely outbreak detection is statistical process control (SPC) charts. Cumulative sum (CUSUM) charts have been especially popular in public health surveillance to detect increases in the rate or counts of disease incidents. However, the use of CUSUM charts in biosurveillance can be challenging when compared to the conventional SPC charts in other application areas. Possible challenges and opportunities include the following issues.

- Since the disease outbreak is a spatiotemporal event, the surveillance method should detect the time when an outbreak occurs as well as the coverage of the outbreak. Fast and accurate detection is important and preferred. Also, the performance of the method should be evaluated with respect to detection time and accuracy.
- Correlation may exist in observed data. This could be in the form of spatial correlation, autocorrelation (over time), or both. A detection method that can take advantage of the correlation, or at least can be applied to the correlated data, is needed.
- The complexity in designing CUSUM charts often relates to the dimensionality of monitoring regions. The control limits of the CUSUM charts are typically calibrated by trial-and-error simulation, but this task can be extremely time-consuming and challenging if the number of monitoring regions is large.
- Surveillance through the CUSUM charts requires intensive modeling of the monitoring system that changes according to the input parameters and the underlying data distribution (which is often unknown).

In this thesis, we focus on spatiotemporal multivariate CUSUM (MCUSUM) charts to detect a disease outbreak. In order to avoid tedious simulations to calibrate the MCUSUM charts, we first propose MCUSUM charts whose control limits

are approximated analytically by assuming the charts are independent; and then we will expand our research to make the charts applicable to any data distribution, even in the presence of the correlation among the charts.

Chapter 2 introduces a stochastic model to investigate realistic strategies that can be used in reaction to emerging outbreaks of pandemic influenza. The model is constructed to represent a typical mid-sized North American city, and calibrated to documented illness attack rates. Reduced attack rates and economic costs due to the intervention strategies are estimated. Chapter 3 reviews MCUSUM charts based on sequential likelihood ratio tests in the presence of spatial correlations, and then proposes an analytical method that approximates the control limits, making the design of the MCUSUM charts much more convenient. We also study how spatial correlation impacts the scheme's outbreak detection performance. Chapter 4 studies another class of MCUSUM chart that can be robust to any underlying distribution of data in the presence of spatial correlation. The chart provides reliable performance by overcoming certain restrictions of the MCUSUM charts developed in Chapter 3. Chapter 5 summarizes the main contributions of the thesis.

Chapter II

REACTIVE STRATEGIES FOR DEVELOPING OUTBREAKS OF PANDEMIC INFLUENZA

In April, 2009, the World Health Organization (WHO) announced the emergence of a new influenza A (H1N1) virus, and on June 11, 2009, it declared that the world was at the start of a new influenza pandemic [68]. WHO reported more than 414,000 laboratory-confirmed cases of H1N1 [67] — a gross underestimate, as many countries simply stopped counting individual cases. The U.S. Centers for Disease Control and Prevention reported widespread influenza activity in forty-six states, with influenza-like illness (ILI) activity in October 2009 higher than what is seen during the peak of many regular flu seasons; and further, “Almost all of the influenza viruses identified . . . are 2009 H1N1 influenza A viruses” [7]. Countries found themselves in the position of having to react to contain already developing Fall outbreaks of influenza due to the new pandemic strain, a position they are likely to find themselves in again if and when future waves of pandemic influenza occur.

Research has suggested that mass vaccination of 60–70% of the population prior to the start of the flu season could effectively contain outbreaks due to pandemic strains [12, 17, 29, 69]; and the public health preparedness plans of most countries have, accordingly, emphasized vaccination intervention strategies. However, the recent experience with H1N1 suggests that high vaccination coverage levels are difficult to achieve. In the case of H1N1, vaccination programs in most northern hemisphere countries started only after the virus was widely circulating. Furthermore, in some countries, supplies of vaccine were limited [9], delivery and administration occurred over a period of several months [37, 42], and there were reports of public skepticism

regarding the necessity and safety of vaccination [36, 38], all of which were strong indicators suggesting that high vaccination coverage would be difficult to achieve. While many institutions in the U.S. and elsewhere strongly encouraged and, in some cases, required workers to be vaccinated against seasonal influenza in 2009, H1N1 vaccination guidelines were focused mostly on people in certain age and high-risk groups [8]. Delays, limited and untimely vaccination supplies, and public reluctance to be vaccinated are likely to reduce the effectiveness of vaccination campaigns [17, 69].

The issues outlined above for the recent outbreak of H1N1 are likely to occur again in future outbreaks of pandemic influenza. In this chapter, we explore the effectiveness of realistic reactive intervention strategies implemented after the beginning of outbreaks of pandemic influenza. We calibrate our model based on data for the H1N1 pandemic (see Tuite et al. [63]), and we investigate the impacts of (i) the moderate vaccination coverage levels which, based on past experience, are likely to be realized, as well as high levels which would be more ideal; (ii) very limited treatment of cases with antivirals and prophylaxis of cases' households with antivirals; and (iii) limited and practical social distancing measures such as five-day closure of individual schools on an as-needed basis, encouragement of liberal leave policies in the workplace, and encouragement of self-isolation. Intervention strategies that combine these approaches are also studied (cf. Halloran et al. [18]). For all intervention strategies, we provide cost estimates associated with morbidity and mortality that take into account direct medical costs as well as economic consequences resulting from school closures and work loss. The research conducted in this chapter has been published in Andradóttir et al. [2].

2.1. Simulation Model

We developed a portable and adaptable stochastic, individual-level simulation model of influenza spread within a structured population. The simulator is similar to models

developed by Longini et al. [29, 30]. A simulated population of 649,565 people was generated stochastically to represent a typical North American city, namely, Hamilton (Ontario), Canada, which was chosen due to availability of demographic and epidemiological data necessary for constructing and calibrating the simulator. Our population is a collection of heterogeneous individuals with various attributes that impact whom they interact with (and hence whom they may infect or get infected by). More specifically, each individual has the following stochastically generated attributes: age, household, playgroup or daycare attended (for pre-school children), school attended (for school-age children), workgroup (for working adults), household census tract and workplace census subdivision, community, and neighborhood. As in Longini et al. [30], a community consists of approximately 2000 people living within the same census tract, and a neighborhood consists of approximately 500 people living within proximity to each other within the same community; also see the recent papers including Aleman et al. [1] and Stroud et al. [57], which incorporate more-detailed individual-level behavior involving larger populations. Age and household-size distributions, shown in Figures 1 and 2, were matched to 2001 Canadian census data [54, 55]. Household census tract assignments were made so that census tract population sizes were consistent with 2006 census statistics [51].

Each individual belongs to three or four contact groups. In particular, each individual belongs to a household, neighborhood, and community. In addition, children younger than 16 belong to either a playgroup, daycare, or school, depending on age; most children in age range 16–18 belong to a school or workgroup; and most adults in age range 19–59 belong to a workgroup. Preschool children were categorized as belonging to a playgroup/daycare, each with 50% probability. We separated secondary schools into middle schools and high schools based on grade to allow different contact group sizes and to make our model more representative of mid-sized U.S. cities. The numbers of playgroups, daycares, elementary, middle, and high schools in each




Figure 1: Age distribution for simulated population




Figure 2: Household size distribution for simulated population

community were based on Longini et al. [30], and were combined with the number of individuals in each category in our simulation population to obtain the contact group sizes.

Workgroups were formed to match 2006 employment statistics [53] as well as census statistics on the geographical distribution of workers [52]. The number of working adults (19–59 years old) was based on census data [52]; and the number of working children (16–18 years old) was based on Ontario data on drop-out rates [5] and the employment rate for ages 15–24 [52]. Rather than representing entire workplace institutions, we formed workgroups of size 20 to represent the typical number of co-workers an individual is likely to have close contact with during the day. Average playgroup, daycare, and lower and upper secondary school (i.e., middle and high school) contact group sizes were chosen following similar reasoning.

2.1.1 Disease Transmission Model

The simulator models influenza transmission over a 180-day period, within the contact groups previously defined. To initiate influenza outbreaks, simulations are seeded with approximately 100 randomly selected initial infectives, with all other individuals considered susceptible (state 0). Susceptible people are assumed to have daily contacts with other individuals in their contact groups, i.e., their household and school or workgroups, as well as with people in their neighborhood and community. Susceptible people have the opportunity, each day, to become infected in their contact groups. The daily probability of infection for each susceptible person is determined by the number of infectious contacts in his contact groups, and on the per-contact probability of transmission for each type of contact. For example, the probability of a susceptible child who attends daycare being infected on a particular day is:

$$1 - \left[\Pr(\text{child is not infected in the household}) \right. \\ \left. \times \Pr(\text{child is not infected in the neighborhood}) \right]$$

$$\begin{aligned} &\times \Pr(\text{child is not infected in the community}) \\ &\times \Pr(\text{child is not infected at the daycare center}) \end{aligned}] .$$

Within each contact group, the probability of infection of a susceptible individual depends on the number of infectious individuals in the group. For example, suppose that k_1 children and k_2 adults in a household are infectious on a particular day. Then the probability of a susceptible household member being infected in that household on that day is:

$$\begin{aligned} &1 - [\Pr(\text{not infected by a child in the household})^{k_1} \\ &\times \Pr(\text{not infected by an adult in the household})^{k_2}] . \end{aligned}$$

The number of infectious people in the contact groups (e.g., k_1 and k_2), are random variables that are updated at the beginning of each day.

Infection of susceptibles depends on the number of infected persons in their contact groups, on the vaccine and antiviral-use status of susceptibles and their infectious contacts, and on age- and contact-group-specific per-contact transmission probabilities (Table 1). The probability that infection is transmitted from an infected person to a susceptible person also depends on whether the infectious person is symptomatic or asymptomatic. Table 1 shows the rates for symptomatic individuals. The transmission rates for asymptomatic individuals are half of those shown in Table 1. These probabilities are based on Longini et al. [29, 30], with adjustments made to calibrate baseline (no intervention) results.

Table 1: Per-contact influenza transmission probabilities within contact groups

Contact Group	Transmission Probability
Household ¹	
Child-to-child	0.8
Child-to-adult	0.3
Adult-to-child	0.3
Adult-to-adult	0.4
Community ²	
Pre-schooler	0.000005
School child	0.000005
Adult (ages 19–52)	0.000075
Adult (ages 53+)	0.000055
Daycares/Playgroups ²	
Daycares	0.028
Playgroups	0.018
Schools ²	
Elementary schools	0.012
Middle schools	0.011
High schools	0.010
Workgroups	0.010

¹ Within households, the probability that a symptomatic child (age 18 years or less) infects a susceptible child is 0.8; that a symptomatic child infects a susceptible adult (at least 19 years old), or that a symptomatic adult infects a susceptible child, is 0.3; that a symptomatic adult infects a susceptible adult is 0.4 [30].

² Probability that a susceptible person in the age or school group is infected through contact with a symptomatic person in the group.

Table 2: Age-group-specific illness attack rates and calibrated attack rates

Age	Simulated illness attack rates by the percentage of adults 53+ years old with pre-existing immunity ¹			Calibrated attack rates (AR)	
	30%	50%	70%	Age	AR
0–4	30.6%	31.0%	30.8%	0–4	29.5%
5–13	53.8%	55.0%	55.2%	5–18	55.9%
14–17	56.0%	57.1%	57.3%		
18–22	48.9%	49.7%	49.7%	19–52	40.8%
23–52	39.6%	39.8%	39.3%		
53–64	21.7%	15.3%	8.8%	53–59	14.3%
65+	19.1%	13.2%	7.5%	60+	11.0%
Overall	36.8%	35.4%	33.5%	Overall	34.1%

¹ See the discussion in Tuite et al. [62].

The simulator is calibrated to match documented illness attack rates and basic reproduction numbers (R_0), defined as the number of cases one case generates on average over the course of its infectious period. Baseline (no-intervention) scenario age-group-specific attack rates were derived using 2009 estimates for the H1N1 basic reproduction number in Ontario [20, 62, 63] (see Table 2). These rates take into account reduced risk in adults born prior to 1957 [20]. A compartmental model parameterized in this way was well-calibrated to observed attack rates during the Fall pandemic wave in Ontario [62]. The simulator’s R_0 value of 1.4 is also consistent with other published reports such as Fraser et al. [13], New Scientist [35] and Yang et al. [69].

Figure 3 depicts a flowchart of the model. The modeled natural history and simulator dynamics parameters, described below and shown in Figure 3, were based on Longini et al. [29, 55]. People infected with influenza first pass through a 1–3 day latent / incubation period (state 1; average length 1.9 days), during which they do not have influenza symptoms. They are not infectious until the last day of the period; at that point, they become half as infectious as if they were to develop

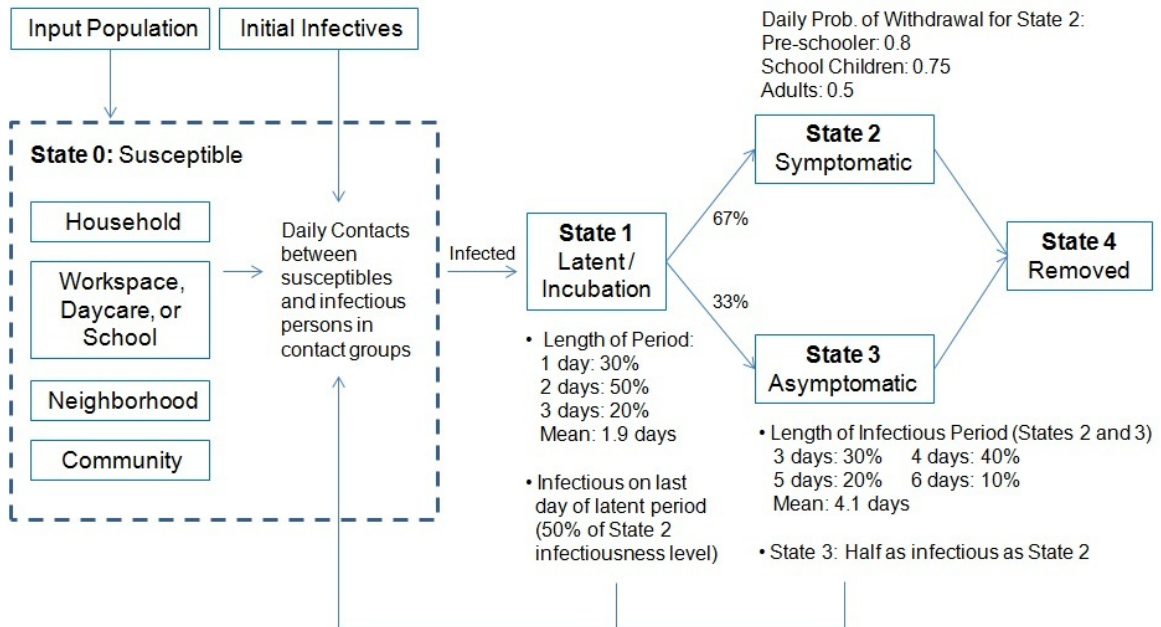


Figure 3: Simulation flowchart and modeled influenza natural history

symptoms in the subsequent period. During that subsequent infectious period, 67% will develop influenza symptoms (state 2), and 33% will be asymptomatic (state 3). Symptomatic infectives are assumed to be twice as infectious as asymptomatics, and have a chance of withdrawing home during each day of illness; upon withdrawal, they only make contacts within their household and neighborhood, with transmission probabilities doubled in the household contact group, until they recover. If a school child withdraws home due to illness, one adult in the household also stays home. Each day in states 2 and 3, an infectious person has a chance to exit the state and be removed from the simulation (i.e., to recover or die — state 4). Probabilities for transition into and out of states are given in Figure 3 and are based on Longini et al. [29, 30].

2.1.2 Intervention Strategies

We modeled a baseline case where no intervention takes place, along with strategies representing various combinations of vaccination, antiviral treatment and household prophylaxis, school closure, and general social distancing (see the results in Tables 3 and 4 and Table 23 provided in the Appendix A). Each component of the strategies is described in detail below. Interventions are triggered in a particular simulation run when the overall illness attack rate reaches 0.01%. Twenty runs of the simulator were performed for each intervention strategy, from which average illness attack rates were calculated. We briefly describe the interventions under consideration.

Vaccination: We model both pre-vaccination as well as reactive strategies, with reactive vaccination programs beginning immediately, 30 days, or 60 days after the trigger. The delays model disruptions in vaccine production and supply chains. We allow enough doses to cover either 35% or 70% of the population. In reactive strategies, we consider cases where (i) all vaccines become available at the same time, and (ii) the doses become available in three equal-sized batches, two weeks apart, due to additional production and supply-chain disruptions. Vaccine efficacy refers to the reduction, after vaccination, in the probability of becoming infected due to contact with an infected person (VE_s), or to the reduction, after vaccination, in the probability of infecting a susceptible contact (VE_i). Vaccine efficacy does not refer to the fraction of individuals having an immunogenic response to the vaccine (which is typically much larger than our measures). We study a low-efficacy single-dose vaccine (efficacy against susceptibility to infection, $VE_s = 0.3$, and efficacy against infectiousness, $VE_i = 0.2$) as well as a moderate-efficacy vaccine ($VE_s = 0.4$, $VE_i = 0.5$) [3]. Each day, our model randomly vaccinates any remaining unvaccinated individuals who are either uninfected or in the latent or asymptomatic phases of infection, all with equal probability based on the number of available doses. Moreover, protection from the vaccine builds over time, with 50% of the vaccine’s efficacy realized upon vaccination,

and full protection after two weeks.

Antiviral Treatment and Household Prophylaxis: We investigate strategies involving treatment of infected individuals with a five-day course of antivirals, as well as strategies that also allow for ten-day prophylaxis of the infected individuals' household members. We assume that 1% of individuals do not complete their course. We use an antiviral efficacy against susceptibility (AVEs) of 0.3 and against infectiousness (AVEi) of 0.7 [30]. Individuals receive direct benefit from antivirals only while they are taking them. Antiviral use is considered alone and in combination with other intervention strategies. It is assumed that antiviral courses are available for 10% of the population and that they are distributed to infected individuals and their household members until the supply is exhausted.

School Closure and Social Distancing: We implement a rolling school closure model, where a daycare or school closes for five days if five or more cases are identified in that group. Given that infected individuals are on average infectious for 4.1 days (see Figure 3), closing schools for fewer than 5 days is unlikely to be very effective. It is possible for these groups to close more than once during the simulation. We also model a reduction in workplace and general community contacts of 20% (i.e., 20% of infected individuals in each contact group will not infect other members of the group). This represents the exercise of a general level of caution, including a modest limitation of contacts within workgroups (e.g., by invoking occasional telecommuting and other self-limiting behaviors, holding fewer large meetings, etc.) and also within the general community (e.g., reduction in attendance in social groups and larger community events, etc.).

Table 3: Attack rates and total costs of interventions with 35% vaccination coverage

Intervention ¹	Delay in initiation of vaccination ²	No post-initiation vaccination delays		Post-initiation vaccination delay ³	
		Attack rate (%)	Cost (US\$m)	Attack rate (%)	Cost (US\$m)
None		34.1	81.1		
A		31.3	75.9		
S		24.0	125.0		
A+S		9.2	48.0		
V _L	Pre-vaccination	26.1	71.1		
V _L	Reactive, no delay	28.8	77.7	28.8	77.7
V _L	30-day delay	29.0	78.1	29.5	79.3
V _L	60-day delay	30.7	82.2	32.2	86.0
V _M	Pre-vaccination	18.8	53.7		
V _M	Reactive, no delay	22.6	62.8	22.8	63.1
V _M	30-day delay	23.0	63.7	24.6	67.5
V _M	60-day delay	27.3	74.1	30.8	82.5
V _L +A	Pre-vaccination	19.3	56.4		
V _L +A	Reactive, no delay	25.2	70.6	25.3	70.8
V _L +A	30-day delay	25.4	71.1	25.7	71.8
V _L +A	60-day delay	26.2	72.9	27.1	75.0
V _M +A	Pre-vaccination	2.1	16.1		
V _M +A	Reactive, no delay	8.1	30.1	10.0	34.3
V _M +A	30-day delay	12.4	40.2	15.8	48.2
V _M +A	60-day delay	18.6	54.7	20.8	60.1
V _L +S	Pre-vaccination	12.7	69.9		
V _L +S	Reactive, no delay	17.3	93.6	17.5	95.7
V _L +S	30-day delay	17.8	96.5	18.3	99.0
V _L +S	60-day delay	18.6	101.9	19.6	108.8
V _M +S	Pre-vaccination	2.3	19.6		
V _M +S	Reactive, no delay	6.8	41.6	8.5	49.4
V _M +S	30-day delay	9.9	56.3	15.4	87.3
V _M +S	60-day delay	13.4	74.7	17.9	95.7
V _L +A+S	Pre-vaccination	1.0	15.9		
V _L +A+S	Reactive, no delay	3.9	29.2	4.5	32.2
V _L +A+S	30-day delay	4.6	32.6	4.9	34.2
V _L +A+S	60-day delay	4.8	33.8	5.4	36.8
V _M +A+S	Pre-vaccination	0.2	11.9		
V _M +A+S	Reactive, no delay	0.5	13.1	0.8	14.9
V _M +A+S	30-day delay	1.2	16.6	1.6	18.6
V _M +A+S	60-day delay	2.0	20.2	2.4	22.0

¹ Abbreviations for modeled interventions: V (vaccination of up to 35% of the population), L (low efficacy), M (moderate efficacy), A (antiviral treatment and household prophylaxis of up to 10% of the population), S (school closure and social distancing).

² Initial supply-chain delays which prevent immediate initiation of vaccination programs after the intervention trigger occurs.

³ Additional supply-chain delays, after initiation of the vaccination program, as a result of which vaccines become available in three equal batches, spaced two weeks apart.

Table 4: Attack rates and total costs of interventions with 70% vaccination coverage

Intervention ¹	Delay in initiation of vaccination ²	No post-initiation vaccination delays		Post-initiation vaccination delay ³	
		Attack rate (%)	Cost (US\$m)	Attack rate (%)	Cost (US\$m)
V _L	Pre-vaccination	12.0	47.0		
V _L	Reactive, no delay	22.2	71.1	22.4	71.6
V _L	30-day delay	22.7	72.4	24.1	75.7
V _L	60-day delay	27.1	83.0	30.4	89.4
V _M	Pre-vaccination	0.2	19.3		
V _M	Reactive, no delay	2.2	25.6	4.6	29.7
V _M	30-day delay	8.1	39.5	13.3	50.2
V _M	60-day delay	22.6	74.0	27.6	83.0
V _L +A	Pre-vaccination	3.3	28.3		
V _L +A	Reactive, no delay	17.3	61.1	17.7	62.0
V _L +A	30-day delay	17.9	62.5	18.4	63.9
V _L +A	60-day delay	19.9	67.4	22.0	72.4
V _M +A	Pre-vaccination	0.1	20.7		
V _M +A	Reactive, no delay	0.6	22.0	1.2	23.3
V _M +A	30-day delay	2.4	26.2	4.4	30.9
V _M +A	60-day delay	6.6	36.1	12.2	49.1
V _L +S	Pre-vaccination	0.7	22.0		
V _L +S	Reactive, no delay	5.9	46.0	7.5	53.1
V _L +S	30-day delay	9.5	63.0	11.0	70.6
V _L +S	60-day delay	13.3	82.6	15.4	96.6
V _M +S	Pre-vaccination	0.0	19.1		
V _M +S	Reactive, no delay	0.2	19.7	0.7	22.0
V _M +S	30-day delay	1.5	25.9	3.2	34.7
V _M +S	60-day delay	6.4	51.2	9.8	69.1
V _L +A+S	Pre-vaccination	0.2	21.3		
V _L +A+S	Reactive, no delay	1.8	28.4	2.6	32.0
V _L +A+S	30-day delay	2.9	33.6	3.2	35.2
V _L +A+S	60-day delay	3.8	37.8	4.6	41.7
V _M +A+S	Pre-vaccination	0.0	20.6		
V _M +A+S	Reactive, no delay	0.1	20.1	0.2	21.6
V _M +A+S	30-day delay	0.5	22.8	0.7	23.8
V _M +A+S	60-day delay	1.2	26.1	1.4	27.4

¹ Abbreviations for modeled interventions: V (vaccination of up to 70% of the population), L (low efficacy), M (moderate efficacy), A (antiviral treatment and household prophylaxis of up to 10% of the population), S (school closure and social distancing).

² Initial supply-chain delays which prevent immediate initiation of vaccination programs after the intervention trigger occurs.

³ Additional supply-chain delays, after initiation of the vaccination program, as a result of which vaccines become available in three equal batches, spaced two weeks apart.

2.2. Economic Cost Estimation

We determine economic costs associated with the influenza outbreaks and modeled intervention strategies. We include medical spending due to illness, costs of antivirals and vaccines, and costs associated with teachers and other working adults staying home due to their own illness, illness of dependent children, or due to school closure. Medical spending includes co-payments and net payments for outpatient visits and hospitalization, as well as prescription and over-the-counter medications for influenza and complications or secondary infections. Costs are stratified by age-group and by low- or high-risk status of individuals with respect to complications of influenza. We also include the present value of earnings lost due to premature mortality.

Cost estimates and probabilities of risk status and of complications and death were taken from Meltzer et al. [33], with costs inflated using 2008 consumer price index and medical price index estimates [6, 23, 32, 58]. These costs are combined with the data on age-specific attack rates, utilized vaccination doses, and days of school closure obtained from our simulation model.

The total cost of each intervention scenario includes the cost of vaccine doses and antiviral courses used, if any; costs associated with parents staying at home with sick children and school teachers, parents, and children staying home due to school closure; costs due to illness-related absence from work; medical costs associated with illness, including outpatient visits, prescription and over-the-counter drugs, and hospitalization; and lost earnings due to death.

We use methods described by Meltzer et al. [33] to quantify most medical and work-loss costs (see also Medlock and Galvani [32]). Table 5 shows the proportions of illnesses assumed to be at high risk for complications among children (0–18 years old), younger adults (19–59 years old) and seniors (over 60). Table 6 shows estimated rates of outpatient visits, hospitalizations, and death used in our calculations for children, adults, and seniors at high risk and not at high risk of complications. We chose the

Table 5: Proportions of influenza cases at high risk for complications¹

Age Group	Proportion at high risk
Children (0–18)	0.064
Adults (19–59)	0.144
Seniors (60+)	0.400

¹ Proportions taken from Meltzer et al. [33], and adapted to our age groups.

Table 6: Outpatient visit, hospitalization, and death rates¹

	Rates per 1000 persons ill		
	Outpatient visits	Hospitalization	Deaths
Not at high risk			
Children	165	0.20	0.014
Adults	40	0.18	0.025
Seniors	45	1.50	0.280
High risk			
Children	289	2.10	0.126
Adults	70	0.83	0.100
Seniors	79	4.00	2.760

¹ Rates taken from Meltzer et al. [33]

‘low’ rate estimates presented in Meltzer et al. [33], which we believe to be most consistent with the relatively low R_0 (1.4) for our model.

Frequency and costs (in US\$) associated with influenza-related outpatient visits, hospitalizations, and deaths are shown in Table 7. All the above costs were combined with age-specific attack rates obtained from our simulation model. In addition, we assume average costs of \$25 per vaccine dose or antiviral course used, consistent with previous reports [48]. Table 8 shows other costs associated with vaccination (i.e., the cost of lost time, travel, and side effects). These costs are based on Bowlby [5], inflated as described above. The vaccination costs are combined with the number of used vaccination doses obtained from our simulation model. We assume that 1% of antiviral users discontinue use due to side effects; medical and other costs associated

with these side effects are not included in our model.

To estimate costs of ill individuals staying home and work-loss associated with parents staying at home with sick children, we multiplied the number of days (obtained from our simulation model) with the inflation-adjusted average value of lost days from Table 7. Similarly, we estimated the average number of teachers at schools and daycares by dividing the total number of such teachers in Hamilton [56] among the schools and daycares in our model. To estimate the cost of lost teacher productivity due to school closures, we multiplied the number of days schools and daycares are closed in our simulation model by the average number of teachers at Hamilton schools and daycares and by the average value of a day of lost work obtained from Table 7.

Table 23 in the Appendix A shows age-stratified and overall illness attack rates for all modeled scenarios, along with total cost estimates. Figure 4 depicts the total cost (US\$m) plotted vs. average overall illness attack rate (%) for each intervention.

Table 7: Outpatient visit, hospitalization, and death costs¹

Outcome category item	Age group		
	Children	Adults	Seniors
Outpatient Visits			
Average no. visits per case	1.52	1.52	1.52
Net payment per visit	\$80.90	\$62.74	\$82.55
Average copayment for outpatient visit	\$8.26	\$6.60	\$6.60
Net payment per prescription	\$41.28	\$59.44	\$59.44
Average prescriptions per visit	0.9	1.8	1.4
Average copayment per prescription	\$4.95	\$4.95	\$4.95
Days lost	3	2	5
Value of 1 day lost	\$91.85	\$141.30	\$91.85
Subtotal: Per-case Outpatient Costs	\$448.86	\$496.50	\$679.47
Hospitalization			
Hospital cost	\$4,847.34	\$9,932.42	\$11,319.26
Net payment per outpatient visit	\$122.17	\$155.19	\$168.40
Average copayment for outpatient visit	\$8.26	\$6.60	\$6.60
Net payment for drug claims	\$42.93	\$69.34	\$67.69
Most likely days lost	5	8	10
Value of 1 day lost	\$91.85	\$141.30	\$91.85
Subtotal: Per-case Hospitalization Costs	\$5,479.92	\$11,293.96	\$12,480.40
Deaths			
Average age (years)	9	35	74
PV earnings lost	\$1,435,750	\$1,466,231	\$93,027
Most likely hospital costs	\$5,671	\$12,555	\$13,718
Subtotal	\$1,441,422	\$1,478,788	\$106,746
Ill but no medical care sought			
Days lost	3	2	5
Value of 1 day lost	\$91.85	\$141.30	\$91.85
Over-the-counter drugs	\$3.30	\$3.30	\$3.30
Subtotal: Per-case ill (no care sought)	\$278.84	\$285.90	\$462.53

¹ Estimates based on figures from Meltzer et al. [33]. Cost estimates inflated by 2008 consumer and medical price indices [6, 23, 58] as appropriate.

Table 8: Costs (in US\$) and impacts of vaccination¹

Item	Probability of side effect	Per-case cost of side effect	Cost scenario (per patient)
Assumed cost of vaccination			\$25.00
Patient time			\$5.65
Patient travel costs			\$5.65
Side effect			
Mild	0.0325	\$94	\$5.04
Guillain-Barre Syndrome (GBS)	0.000002	\$100,800	\$0.33
Anaphylaxis	0.000000157	\$2,490	\$0.0006

¹ Estimates based on figures from Meltzer et al. [33]. Travel and side effect cost estimates inflated by 2008 consumer and medical price indices [6, 23, 58] as appropriate.

2.3. Simulation Results

With no intervention, the average overall illness attack rate is 34.1%, with an estimated total cost of \$81.1 million (Table 3). Pre-vaccination of 35% of the population with a low-efficacy vaccine reduces the average overall illness attack rate to 26.1% (total cost \$71.1 million), and with a moderate-efficacy vaccine to 18.8% (total cost \$53.7 million). Not surprisingly, pre-vaccination of 70% of the population is more effective (overall average illness attack rate 12.0%, total cost \$47.0 million for a low-efficacy vaccine; and 0.2% and \$19.3 million with a moderate-efficacy vaccine; see Table 4).

Reactive vaccination alone, of 35% of the population with a low-efficacy vaccine delivered in three batches, reduces the overall average illness attack rate to 28.8% (or 22.8% with a moderate-efficacy vaccine), with a total cost of \$77.7 million (\$63.1 million with a moderate-efficacy vaccine). Thirty- and 60-day delays in initiation of reactive vaccination, with vaccines delivered in three batches, result in attack rates of 29.5% (total cost \$79.3 million) and 32.2% (total cost \$86.0 million), respectively, for a low-efficacy vaccine, and 24.6% (total cost \$67.5 million) and 30.8% (total cost \$82.5 million), respectively, for a moderate-efficacy vaccine. Figure 5 shows daily

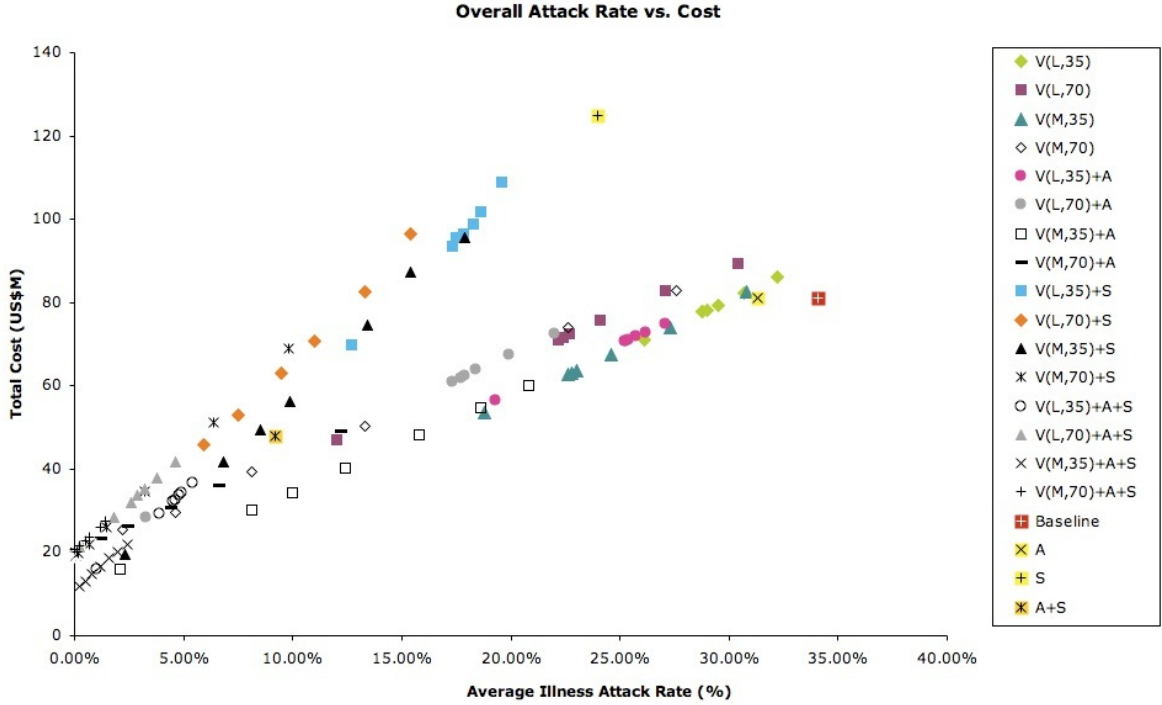


Figure 4: Total cost of intervention strategies vs. the average illness attack rate

attack rates for (i) the case of 70% coverage of low-efficacy vaccine with 60-day initial delay, and (ii) the baseline case. For case (i), the vaccine is given on the 60th day followed by receipt of vaccine after two additional two-week delays (see arrows in Figure 5). Note that vaccine given on the 60th day decreases the attack rate compared to the baseline; but the two subsequent receipts of vaccine do not result in additional benefits. Clearly, with a 60-day delay, interventions occur too late in the epidemic to have any meaningful effect (see Figure 5).

Antiviral use at low (10%) coverage alone results in an overall attack rate of 31.3% (total cost \$75.9 million). School closure and social distancing alone result in an attack rate of 24.0%, with a total cost of \$125.0 million.

Suppose we combine reactive low-efficacy vaccination of 35% of the population delivered in three batches, antivirals (10% coverage), and school closure and social distancing. Then the overall average illness attack rate is 4.5% (total cost \$32.2




Figure 5: Daily attack rates

million) if no delays occur in the initiation of vaccination, and 5.4% (total cost \$36.8 million) if a 60-day delay occurs. With a moderate-efficacy vaccine, the attack rate for this last scenario reduces to 2.4% (total cost \$22.0 million). Similar relationships between interventions are apparent for interventions with 70% vaccination coverage, shown in Table 4. Vaccination coverage of 70% with a moderate-efficacy vaccine, combined with antiviral treatment and school closure, is highly effective, even with an initial 60-day delay and additional supply-chain disruptions (average illness attack rate 1.4%, total cost \$27.4 million).

We note that the results when all vaccines are available at the same time are better than those involving delivery in batches, and sometimes significantly so, especially for a moderate-efficacy vaccine (Tables 3 and 4). Figures 6 and 7 illustrate the comparative illness attack rates of the various intervention strategies discussed above for all combinations of low/moderate-efficacy vaccine delivered in three batches and at 35%/70% coverage as a function of the initial delay in vaccination implementation due to supply-chain disruptions. The impact of vaccinating 70% of the population,

rather than 35%, ranges from moderate to substantial, with the increased coverage being most beneficial when the vaccine is delivered in a timely manner, and the vaccine is either of moderate efficacy or of low efficacy applied in combination with other intervention strategies.

Complete (age-stratified and overall) average illness attack results for all modeled interventions are given in Table 23 in Appendix A. The comparative effectiveness of interventions is similar when age-group-specific results are studied.

Figure 8A illustrates attack rate and total cost combinations for interventions that result in at least a 75% reduction in cost compared to no intervention. Abbreviations for modeled interventions are PV (pre-vaccination), V (vaccination), L (low-efficacy), M (moderate efficacy), 35 (35% coverage of population), 70 (70% coverage), A (antiviral treatment and household prophylaxis of up to 10% of the population), and S (school closure and social distancing). Multiple occurrences of each plotting symbol may occur; occurrences at higher costs and illness attack rates represent interventions with longer supply-chain delays. In Figure 8A, the closer to the origin, the more desirable an intervention is in terms of total cost and average illness attack rate. Aside from pre-vaccination strategies, we see that 70% reactive vaccination with a moderate-efficacy vaccine and school closure and social distancing, or even 35% reactive vaccination with a moderate-efficacy vaccine, antiviral use, and school closure, also result in substantial reductions in cost and attack rates. Figure 8B illustrates attack rate and cost results for interventions that result in more-modest 50–75% reductions in cost compared to no intervention. Once again, several strategies combining vaccination, antiviral use, and school closure/social distancing are competitive with pre-vaccination.

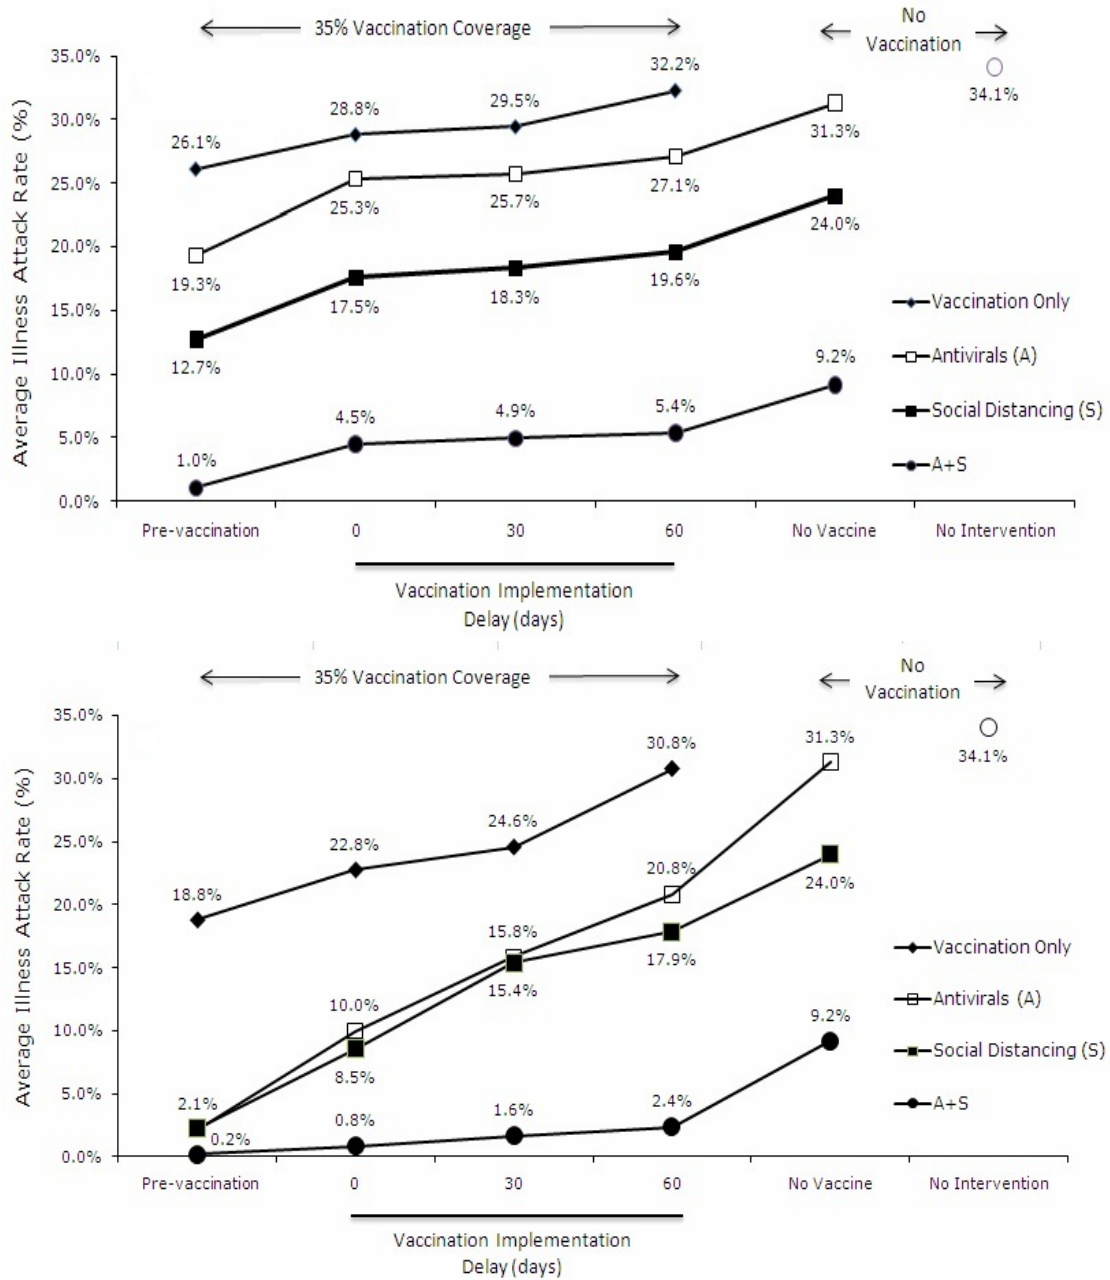


Figure 6: Illness attack rates for modeled interventions with 35% vaccine coverage

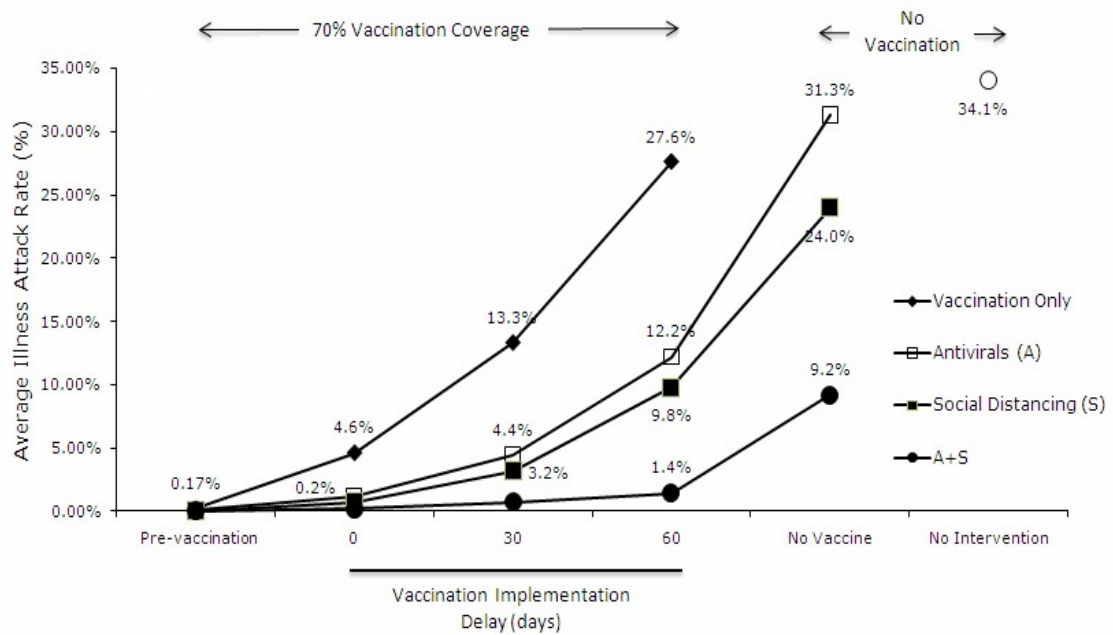



Figure 7: Illness attack rates for modeled interventions with 70% vaccine coverage

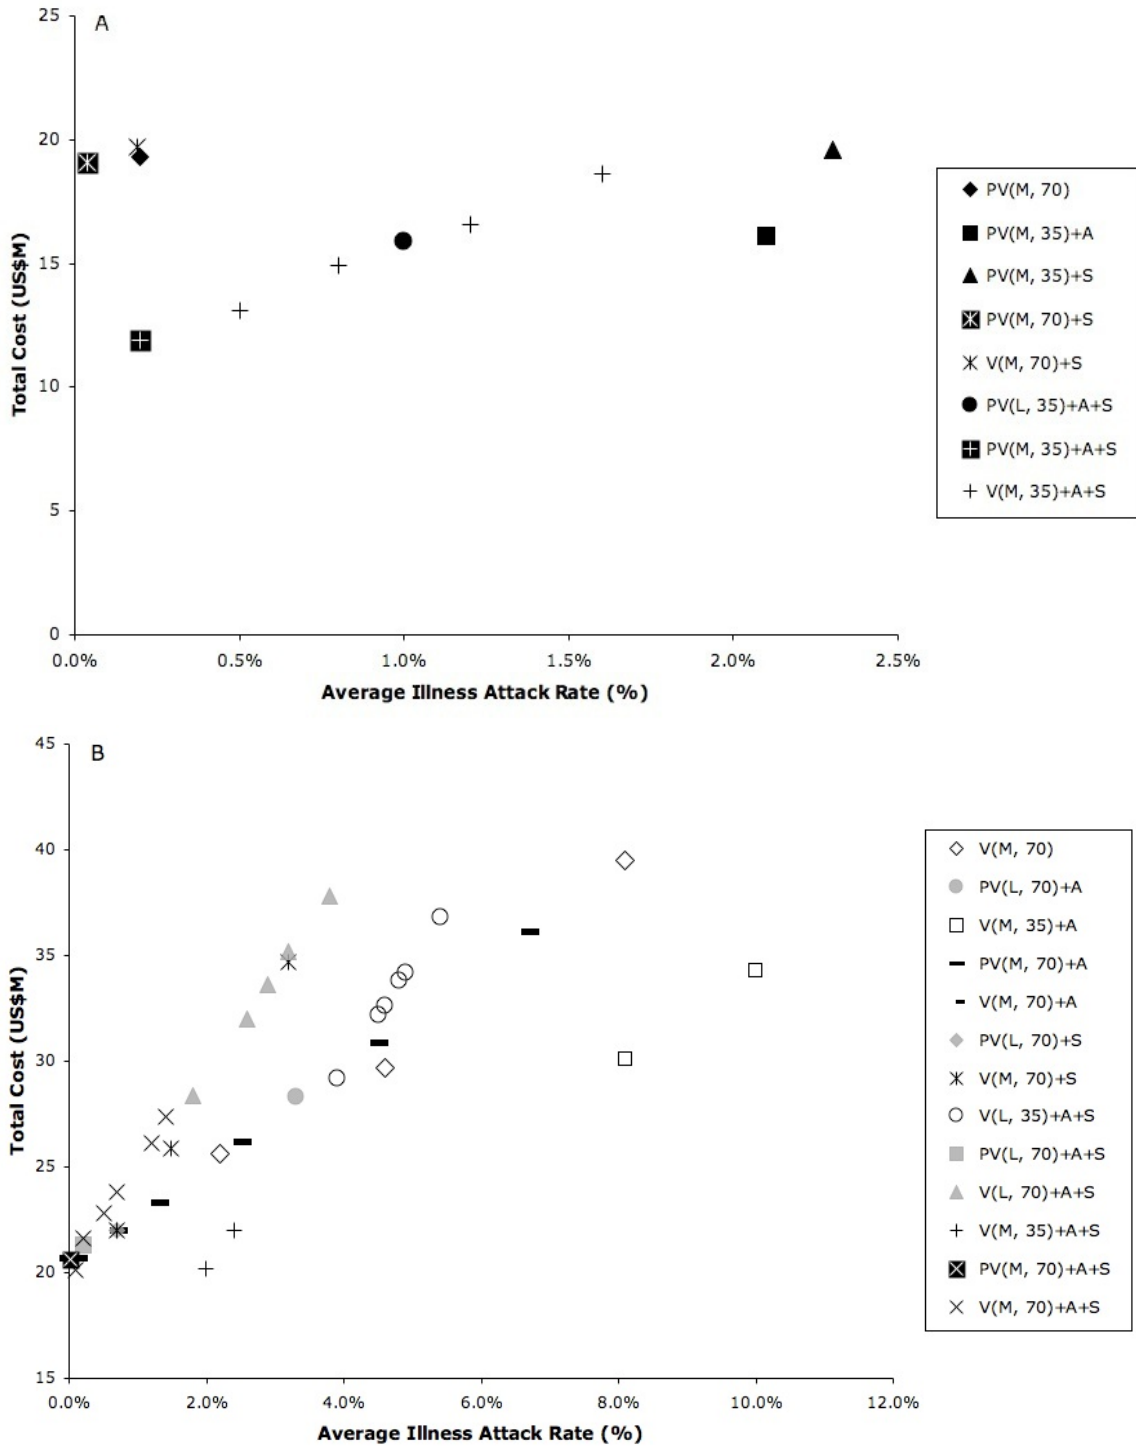


Figure 8: Total cost of modeled intervention strategies vs. average illness attack rate

2.4. Discussion

Previously published research has shown that pre-vaccination of 60–70% of the population can contain seasonal as well as pandemic influenza, but that delays in vaccination can greatly reduce the effectiveness of the vaccination programs [5–8]. Our model confirms these results for moderate-efficacy vaccines (Tables 3, 4, and 23). However, vaccination efforts in countries such as the U.S., Canada, and others began well after the first waves of H1N1 activity, and it is reasonable to believe that the same will be true in future outbreaks of pandemic influenza. In particular, in the event of an outbreak, it will likely take time to achieve high levels of vaccination coverage, and, if past experience with seasonal influenza vaccination campaigns is an indication, it is plausible that only low or moderate coverage will eventually be achieved. The results of our simulation model show that delayed and low-coverage reactive vaccination strategies (with a low-efficacy vaccine, plus limited use of antivirals) will not be enough to mitigate the pandemic or to significantly reduce total costs associated with influenza morbidity and mortality (based on results from Table 3, average illness attack rates are only reduced by 26% and total costs by 13%, compared to no intervention).

According to our model, combining rolling, limited-duration, as-needed closures of individual schools and a practical social distancing policy with 35% reactive low-efficacy vaccination coverage and low-level (10%) antiviral use can reduce illness attack rates by 89% compared to no intervention, as well as total costs by 64%. Similarly, combining interventions in this manner reduces overall attack rates by 99% and costs by 84% when a moderate-efficacy vaccine is available. This strategy remains highly effective even when delays in implementing vaccination of up to 60 days occur. Previously published results have left open the question of how costly interventions involving school closure might be [17]. Our results show that reactive combination strategies that include practical school closure measures, when diligently

implemented, can reduce total costs associated with influenza morbidity and mortality substantially.

Our model has several limitations. We do not consider vaccination strategies targeted to high-risk groups, which could reduce costs associated with complications from influenza. We have not modeled co-circulating strains of seasonal and pandemic influenza or possible resistance to antiviral drugs (although, to mitigate this limitation, our model assumes only low coverage with antivirals, as well as interventions without antivirals). As is always the case with simulation models, continuing follow-up analyses are needed, including: (i) sensitivity to model parameters; (ii) sensitivity to model intervention triggers (e.g., overall illness attack rate, numbers of cases detected in schools, etc.); (iii) sensitivity to R_0 , which can be heterogeneous across cities and countries; and (iv) results for new H1N1 natural history and transmission parameters, and new cost estimates for complications resulting from H1N1 illness, as they become known.

Our model has several strengths. We model a large, realistic, heterogeneous population, base the simulation model on well-studied and documented stochastic simulators, calibrate to actual H1N1 attack rates and most-likely R_0 values, and have the ability to model large numbers of scenarios in a relatively short amount of time on a desktop platform. The model also provides cost estimates that are useful for making policy decisions about potentially expensive interventions. In particular, we model and analyze a variety of interventions and combinations of interventions in terms of costs and efficacy. We also take into consideration reactive strategies incorporating supply-chain delays, and we identify strategies that effectively contain outbreaks and costs even in the presence of supply-chain delays, low vaccine efficacy, and low vaccine coverage.

Our model illustrates the epidemiological effectiveness of a combination strategy

involving short-term closures of individual schools on an as-needed basis, other practical social distancing activities, reactive vaccination of 35% or more of the population, and limited use of antivirals for treatment and prophylaxis. The model also quantifies the cost savings for this and alternative reactive strategies. Public health authorities should consider placing renewed emphasis on such combination strategies when reacting to possible additional waves of the recent pandemic, or to new waves of future pandemics.

Chapter III

SPATIOTEMPORAL BIOSURVEILLANCE: CONTROL LIMIT APPROXIMATION AND THE IMPACT OF SPATIAL CORRELATION

Control charts are used to determine whether or not a process being monitored is in a state of statistical control. A charted process typically runs until it exceeds certain control limits, indicating potential problems that may require corrective actions. The idea is for the chart to quickly detect an out-of-control state of the process, while simultaneously avoiding false positives (declaring problems when they do not actually exist). To this end, we denote by ARL_0 the expected number of samples (average run length) until a false alarm occurs when a monitored process is actually in control. An out-of-control average run length is denoted as ARL_1 , and it represents the expected number of samples until an alarm when an out-of-control process is monitored. ARL_1 s are often used as performance measures for timely detection of an out-of-control process. Generally speaking, procedures having low ARL_1 s while simultaneously possessing a pre-specified ARL_0 are desirable.

Once control limits are determined given a particular ARL_0 , statistics constructed from observations of a monitored process are plotted over time. As long as the statistics fall within the control limits, the process is assumed to be in control and no corrective action is necessary. However, a point that plots outside of the control limits is interpreted as evidence of the out-of-control state, and subsequent analysis of the control chart can help find the sources of variation, which can then be adjusted to bring the process back into control. Shewhart-type charts use either raw observations or averages of a few observations as monitoring statistics, while CUSUM-type charts

use cumulative sums of observations [34]. There are other types of control charts that incorporate different forms of statistics; for example, see Lucas and Saccucci [31] and Woodall et al. [66]. Our interest in this chapter is on CUSUM-type charts.

A CUSUM control chart is a sequential analysis technique suggested by Page [39] to monitor and detect changes in the parameter of interest. As its name implies, a CUSUM chart raises an alarm if the cumulative sum statistic exceeds pre-specified control limits. Since a CUSUM control chart accumulates information given during the entire monitored time period, it is relatively sensitive to small process shifts. Due to this proclivity, CUSUM charts have been found to be useful in the area of bio-surveillance for detecting a slowly emerging disease cluster, which usually does not result in a sudden large shift in the number of disease counts. Lawson [27], Sonesson and Bock [50], Tsui et al. [60] and Woodall [64] provide reviews of many statistical surveillance methods, including CUSUM charts, in public health. Most work using CUSUM charts assumes that the underlying observations (i.e., the observed disease counts) follow Poisson or normal distributions, at least approximately. Either sufficiently large numbers of counts or some transformation to normality might be necessary to use the normality assumption on the counts.

In biosurveillance, an observation (disease count) is obtained from a region every time unit. If one monitors only a single region to detect when a disease outbreak occurs there, and if the collected observations from the single region are univariate, then a univariate CUSUM chart [39, 40] is used for temporal surveillance of the region. In the context of temporal surveillance, Fricker et al. [14] compare the performance of a CUSUM chart against two Shewhart-type charts and one weighted moving average chart [22, 70]. They find that the CUSUM chart performs significantly better than the other charts across all scenarios they evaluated. When the observations follow a Poisson distribution, Han et al. [19] compare the performance of three detection methods: A scan statistic chart based on a fixed number of most-recent observations, a

CUSUM chart, and an exponentially weighted moving average (EWMA) chart. Their simulation study shows that the CUSUM chart and the EWMA chart outperform the scan statistic chart.

When the decision maker monitors a number of adjacent regions at the same time, we now need a spatiotemporal surveillance method to detect when and in which regions an outbreak occurs. At a given time, each region generates an observation and we monitor a multidimensional vector formed from these observations, which requires a multivariate CUSUM (MCUSUM) chart [10, 21, 41]. Several researchers present various MCUSUM charts for spatiotemporal surveillance. Rogerson [43] introduces an MCUSUM chart using a statistic from Tango [59] which can describe the spatial disease pattern in the entire study region. Rogerson and Yamada [46] compare two multivariate control chart schemes: running multiple univariate CUSUM charts (one for each region separately) and running one MCUSUM control chart.

Multiple univariate CUSUM charts [65] are well suited to the situation where there is little or no spatial correlation, and where changes are anticipated in a relatively small number of regions. On the other hand, if the spatial correlation is strong and can be well estimated, then the MCUSUM chart may be preferred, raising preferred, raising alarms faster. This is because statistics in each univariate CUSUM chart are constructed from observations of a single region, and thus interaction among regions due to spatial correlation are not reflected in the statistics. In biosurveillance, one observation from a region often correlates to observations from surrounding regions due to spatial correlation, and hence an outbreak is likely to develop in a cluster, i.e., a group of neighboring regions. With this in mind, Rogerson and Yamada [45] develop a CUSUM chart that uses local statistics. A local statistic is defined as a weighted sum of observations in a local cluster that is used to detect the outbreak cluster, where the weights can decline with increasing distance from the center of the local cluster. Sonesson [49] defines a spatial cluster as a group of regions in a circle

with a varying radius. However, spatial correlation among regions is not considered. Following a similar definition of a cluster from Sonesson [49], Jiang et al. [24] relate existing correlations among regions are present.

Although Jiang et al. [24] propose the promising MCUSUM chart for monitoring many regions simultaneously while considering correlations among the regions, the control limit calibration for the chart primarily relies on trial-and-error simulation. This becomes inconvenient and time-consuming as the number of regions and the number of clusters become large. In this chapter, we introduce a method that analytically calculates control limits of the MCUSUM chart from Jiang et al. [24] and its variation which we call the separated MCUSUM (SMCUSUM) chart. We also investigate the practical range of the scan radius for the control limit approximation method.

Most of the spatiotemporal surveillance literature considers detection delay as a main performance measure of interest, e.g., an ARL_1 or a conditional expected delay [50]. However, the detection delay alone is not a sufficient measure to evaluate the performance of surveillance methods designed to answer both when and where the outbreak occurs. Accurate detection of geographical locations of outbreak regions is as important as prompt detection. Hence, spatial identification accuracy needs to be measured as well as detection delay. A few papers including Jiang et al. [24] and Rogerson and Yamada [46] consider spatial correlation among regions with respect to detection delay, but have not thoroughly studied the impact of spatial correlation on spatial identification accuracy. In this chapter, we discuss detection delay and spatial identification accuracy as detection performance measures, and we investigate the impact of the outbreak radius and spatial correlation on the two performance measures.

The chapter is organized as follows. In Section 3.1, we review MCUSUM charts that use the concept of spatial clusters in biosurveillance. In Section 3.2, we propose

MCUSUM charts with an ARL approximation method, which can be used to set control limits without trial-and-error simulation. Section 3.3 studies the practical range of the scan radius and the performance of the approximation method. Section 3.4 discusses the performance of the MCUSUM chart under various experimental configurations, in addition to the impact of the outbreak radius and spatial correlation on detection performance. A summary and conclusions follow in Section 3.5. The research conducted in this chapter is in print at IIE Transactions [28].

3.1. Background

In this section, we define notation, introduce our problem, and review an MCUSUM chart for spatiotemporal biosurveillance that uses statistics constructed from spatial clusters.

Suppose that $p = M \times N$ regions in a rectangular shape are considered for monitoring, and $q_c = (m_c, n_c)$, where $m_c = 1, \dots, M$ and $n_c = 1, \dots, N$, represents the two-dimensional space coordinate of region c in a set $\mathcal{P} = \{1, 2, \dots, p\}$ as shown in Figure 9. At each time t , a $p \times 1$ observation vector $\mathbf{y}_t = (y_{t1}, y_{t2}, \dots, y_{tp})'$ is monitored. If the baseline mean value in the in-control state, θ_c , and the marginal variance, σ_c^2 , of each component y_{tc} are known for all $c \in \mathcal{P}$, a standardized observation vector $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$ is obtained by setting $x_{tc} = (y_{tc} - \theta_c)/\sigma_c$. The quantity \mathbf{x}_t is assumed to follow a multivariate normal distribution having probability density function $f_{\boldsymbol{\mu}}(\mathbf{x})$ with a mean vector $\boldsymbol{\mu}$ and known variance-covariance matrix Σ , which is constant over times $t = 1, 2, \dots$. Since the observation vectors are standardized, the variance-covariance matrix and the correlation matrix are interchangeable.

Each component of \mathbf{x}_t corresponds to a standardized observation from each spatial region at time t , so an MCUSUM chart on \mathbf{x}_t (or a subset of \mathbf{x}_t) can be constructed for spatiotemporal biosurveillance [60, 61]. MCUSUM charts are developed to detect a shift as soon as possible after an unknown change time ν when the null parameter

c						
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	32	33	34	35
36	37	38	39	40	41	42
43	44	45	46	47	48	49

$q_c = (m_c, n_c)$						
(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)	(1, 7)
(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)	(2, 7)
(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)	(3, 7)
(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)	(4, 7)
(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)	(5, 7)
(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)	(6, 7)
(7, 1)	(7, 2)	(7, 3)	(7, 4)	(7, 5)	(7, 6)	(7, 7)

Figure 9: Coordinate expression of regions ($p = 7 \times 7$)

$\boldsymbol{\mu}_0 = \mathbf{0}$ is shifted to the alternative parameter $\boldsymbol{\mu}_1$, which indicates an out-of-control state.

Under the assumption that the shape of the outbreak coverage is a circle, Jiang et al. [24] define possible spatial clusters of outbreak regions as $O^{c,r} \equiv \{j \mid \text{dist}(q_j, q_c) \leq r, j \in \mathcal{P}\}$ with scan radius r from region c , where $\text{dist}(a, b)$ denotes the Euclidean distance between a and b . Since the center of the circle-shaped coverage is located in c , hereafter, we call c the center region. Note that the scan radius r can be either constant or variable. When r is variable, r is set to be in a range from $[0, \sqrt{M^2 + N^2}]$ or bounded by a given upper limit r_u . However, since the number of regions in $O^{c,r}$ changes only at certain values of r such as 0, 1, $\sqrt{2}$, 2, and so on, it is good enough for practical purposes to consider a finite number of possible values for r . Let \mathcal{R} denote the set of possible settings for r , i.e., $\mathcal{R} = \{r_1, r_2, \dots, r_u\}$, where u is the cardinality of \mathcal{R} .

When a possible outbreak cluster $O^{c,r}$ is scanned, we test if the mean levels of the regions in $O^{c,r}$ are shifted, while the mean levels of the other regions are the same as for the in-control state. Different shift vectors $\boldsymbol{\mu}_{c,r}$ are considered for different $O^{c,r}$ choices, and a set of possible shift vectors can be defined as $\{\boldsymbol{\mu}_{c,r} \mid c \in \mathcal{P}, r \in \mathcal{R}\}$.

Then, a CUSUM chart is designed to detect a shift from $\boldsymbol{\mu}_0$ to $\boldsymbol{\mu}_{c,r}$.

Jiang et al. [24] introduce the following monitoring statistic,

$$S_t^{**} \equiv \max_{c,r} S_t^{c,r} \equiv \max_{c,r} \max_{1 \leq \nu^* \leq t} \sum_{i=\nu^*}^t \ell_i^{c,r}, \quad t = 1, 2, \dots,$$

where

$$\ell_i^{c,r} \equiv \log \frac{f_{\boldsymbol{\mu}_{c,r}}(\mathbf{x}_i)}{f_{\boldsymbol{\mu}_0}(\mathbf{x}_i)} = \boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \left(\mathbf{x}_i - \frac{\boldsymbol{\mu}_{c,r}}{2} \right), \quad i = 1, 2, \dots, t.$$

When a homogeneous shift magnitude $\delta > 0$ is assumed over all outbreak regions, the j th component of $\boldsymbol{\mu}_{c,r}$ is δ if $j \in O^{c,r}$ and 0 otherwise. An alarm is signaled as soon as $S_t^{**} > h^{**}$, where h^{**} is a control limit that determines the operating characteristics of the monitoring chart. When a signal is raised for the first time, we record (i) the current time t as the time when the control chart detects the shift and (ii) ν^* as the estimated change time though the true change time ν is still unknown. It is notable that $S_t^{c,r}$ itself is an MCUSUM statistic that can be calculated recursively by

$$S_t^{c,r} = \max(0, S_{t-1}^{c,r} + \ell_t^{c,r}), \quad t = 1, 2, \dots,$$

where $S_0^{c,r} = 0$ for all $c \in \mathcal{P}$ and $r \in \mathcal{R}$. Therefore, one MCUSUM statistic is recorded for each spatial cluster, and the maximum of all pu of the MCUSUM statistics is used as the monitoring statistic S_t^{**} . Hereafter, unless otherwise specified, we call the control chart monitoring S_t^{**} the JMCUSUM chart.

Jiang et al. [24] test their JMCUSUM chart under various configurations of outbreak coverage, shift magnitude, correlations among regions, and scan radius. For a given shift, the best (smallest) ARL_1 is observed when the scan radius r matches the radius of the actual outbreak coverage, r_{out} . Therefore, using the correct scan radius is important for detecting an outbreak cluster quickly, and the use of a fixed scan radius is recommended if the outbreak radius is known. If no information is available for the outbreak radius, using a variable scan radius seems better than using a fixed scan radius whose value can be quite different than the true outbreak radius. When

using the variable scan radius, the range of the scan radius should be chosen carefully — to be wide enough to include the actual outbreak radius, but not too wide to hinder the detection power of the analysis.

One of the critical decisions that must be made in designing a control chart is that of specifying a control limit, e.g., the control limit h^{**} for S_t^{**} . In order to find the appropriate control limit that yields a target ARL_0 , researchers often rely on tedious trial-and-error simulation, as do Jiang et al. [24]. These simulations can be very time-consuming, especially in the JMCUSUM chart that uses statistics constructed from spatial clusters with different sizes, because the computational burden greatly increases as the numbers of monitoring regions and radius settings increase. If the control limits can be found quickly, we can save a great deal of time in setting up the JMCUSUM chart for spatiotemporal biosurveillance.

3.2. Design of the JMCUSUM Chart

In this section, we use an analytical approximation to search for the control limits of an MCUSUM chart.

By expanding a formula which approximates ARLs of a CUSUM chart in Kim et al. [25], we derive the following result for the MCUSUM statistic $S_t^{c,r}$:

$$ARL \approx \begin{cases} \frac{\Omega^2}{2d^2} \left\{ \exp\left[-\frac{2d(H+1.166\Omega)}{\Omega^2}\right] - 1 + \frac{2d(H+1.166\Omega)}{\Omega^2} \right\}, & \text{if } d \neq 0, \\ \left(\frac{H+1.166\Omega}{\Omega}\right)^2, & \text{if } d = 0, \end{cases} \quad (3.2.1)$$

where $d \equiv E[\ell_t^{c,r}]$, $\Omega^2 = \text{Var}[\boldsymbol{\mu}'_{c,r}\boldsymbol{\Sigma}^{-1}\boldsymbol{x}_t]$, and H is the control limit of the MCUSUM statistic $S_t^{c,r}$. Note that H is not necessarily the same as h^{**} . We can analytically calculate control limits, and even ARL_1 s of S_t^{**} in some special cases, using Equation (3.2.1). Since the solution of Equation (3.2.1) does not exist in a closed form, one may refer to Rogerson [44] for a direct approximation or use a numerical approach. We take the latter approach to solve the equation.

3.2.1 Fixed Scan Radius

If r_{out} is known, Jiang et al. [24] recommend fixing $r = r_{\text{out}}$ for good ARL_1 performance of the JMCUSUM chart. When r is fixed, our monitoring statistic becomes

$$S_t^{**} = \max_c S_t^{c,r}, \quad t = 1, 2, \dots$$

This means that the number of possible spatial clusters $O^{c,r}$ is reduced to p , and one MCUSUM statistic, $S_t^{c,r}$, is monitored for each spatial cluster. An alarm is raised when any of the p MCUSUM statistics exceed a common control limit h^{**} . If an alarm is incurred by $S_t^{c,r}$, then c is considered to be the center of the outbreak.

When an ARL_0 for $S_t^{c,r}$ is specified and the monitored process is in control, Equation (3.2.1) easily enables us to approximate H for $S_t^{c,r}$. Therefore, if we monitor only one spatial cluster $O^{c^*,r}$, then $S_t^{**} = S_t^{c^*,r}$, and we can set $h^{**} = H$ directly. However, finding h^{**} yielding a target ARL_0 for S_t^{**} is no longer a simple matter if we monitor more than one cluster.

Suppose we monitor $p > 1$ spatial clusters at the same time. Then the run length of S_t^{**} becomes equivalent to the minimum run length of p MCUSUM statistics, and this implies that an ARL_0 for each $S_t^{c,r}$ should be set to a much larger value than the targeted ARL_0 for S_t^{**} . When we assume independence among the $S_t^{c,r}$ for $c \in \mathcal{P}$, we can expect that the run length of $S_t^{c,r}$ is approximately p times longer than the run length of S_t^{**} . For this reason, we can obtain h^{**} by calculating the H value that yields an average run length of $p\text{ARL}_0$ for $S_t^{c,r}$. This ARL_0 adjustment is similar to a Bonferroni adjustment mentioned in David [11] and Rogerson and Yamada [46]. Note that the $S_t^{c,r}$ s may not be independent and the approximation quality can be diminished if $r > 0$ or positive spatial correlation exists. In any case, the JMCUSUM chart for the fixed scan radius, denoted as the JMCUSUM-F chart, can be implemented as follows.

The JMCUSUM-F chart

1. Pick any spatial cluster $O^{c,r}$. Then
 - (a) Solve Equation (3.2.1) for H using the corresponding $S_t^{c,r}$ with $p\text{ARL}_0$.
 - (b) Set $h^{**} = H$.
2. Raise an alarm at t if $S_t^{**} > h^{**}$.

After h^{**} is obtained, Equation (3.2.1) can be used again to approximate the ARL_1 of the JMCUSUM-F chart, denoted as ARL_1^{JF} . In the out-of-control state, \mathbf{x}_t is assumed to follow a multivariate normal distribution with parameters $\boldsymbol{\mu}_1$ and Σ . The quantity ARL_1^{JF} can be approximated by the minimum of the ARL_1 s calculated for $c \in \mathcal{P}$ with the shift vector $\boldsymbol{\mu}_1$, which is likely to be achieved at the outbreak center. In Section 3.4, we will confirm that most of the alarms are truly raised by one spatial cluster whose center matches the outbreak center; and thus the ARL_1 approximation method based on that cluster is acceptable when r is fixed.

3.2.2 Variable Scan Radius

If r_{out} is unknown, the use of a variable r is recommended. Let us denote the JMCUSUM chart with the variable scan radius as the JMCUSUM-V chart. In this case, S_t^{**} becomes equivalent to the maximum of pu MCUSUM statistics, and an alarm is raised as soon as any of the pu MCUSUM statistics hits a common control limit h^{**} calibrated by trial-and-error simulation. If an alarm is incurred by $S_t^{c,r}$, then c is considered as the center of the outbreak and r is interpreted as the radius of the detected outbreak cluster.

Since the number of nonzero components in $\boldsymbol{\mu}_{c,r}$ depends on r , Equation (3.2.1) provides different control limits for different r settings. For this reason, unlike the JMCUSUM-F chart, it is not easy to analytically approximate a common control limit h^{**} for the JMCUSUM-V chart. Instead, we suggest a new MCUSUM chart

that does not require the common control limit h^{**} .

From Table I in Jiang et al. [24], we notice the empirical fact that the ARL_1 performance of the JMCUSUM-V chart seems to be close to the minimum of the ARL_1 s of various JMCUSUM-F charts, each applied with $r = r_1, r = r_2, \dots$, and $r = r_u$, respectively. The minimum is usually achieved when the center of a spatial cluster matches the outbreak center and $r = r_{out}$ no matter what r_{out} is tested. This can be interpreted to mean that the JMCUSUM-V chart behaves like a mixture of the multiple JMCUSUM-F charts. Based on this intuition, we suggest a new MCUSUM chart named the separated-MCUSUM (SMCUSUM) which develops an MCUSUM control chart for each scan radius $r \in \mathcal{R}$ separately:

$$S_t^{*r} \equiv \max_c S_t^{c,r}, \quad t = 1, 2, \dots$$

Then the SMCUSUM chart is implemented as follows.

The SMCUSUM chart

1. For each $r \in \mathcal{R}$,
 - (a) Pick any spatial cluster $O^{c,r}$.
 - (b) Solve Equation (3.2.1) for H using the corresponding $S_t^{c,r}$ with $puARL_0$.
 - (c) Set $h^{*r} = H$.
2. Raise an alarm at t if any $S_t^{*r} \geq h^{*r}$ for $r \in \mathcal{R}$.

Note that the JMCUSUM-F chart with an analytical control limit can be a special case of the SMCUSUM chart. In Section 3.3, we will show by experimentation that the ARL_0 s of the SMCUSUM chart are indeed close to the target ARL_0 . The good news is that, by using the SMCUSUM chart, we do not need to search for h^{**} via trial-and-error simulation. This advantage saves significant time, especially when we deal with large problems with numerous possible values of the scan radius r .

3.3. Approximation Accuracy

In this section, we test the accuracy of the control limit and ARL_1 approximations under various experimental configurations.

3.3.1 Experimental Setup

We construct an archetypal example with 7×7 and 20×20 regions. The target ARL_0 is set to 100 and 370, respectively. The range of the scan radius r is discussed in the next subsection. For the out-of-control state, we assume a homogeneous shift magnitude with $\delta = 1$, and consider outbreak clusters with various r_{out} values when the outbreak center c_{out} is located at the center of the 7×7 regional system (region 25). Figure 10 illustrates outbreak regions with $c_{\text{out}} = 25$ colored in gray for $r_{\text{out}} = 0, 1, 2$.

We consider two different correlation structures. The first correlation structure Σ_1 denotes the $p \times p$ correlation matrix given in Jiang et al. [24], where any pair of adjacent regions has a correlation coefficient $0 \leq \rho \leq 1$, any two cross-adjacent regions have correlation coefficient $\rho/2$, and any other pairs of regions have no correlation. For the second correlation structure Σ_2 , any pair of regions has a correlation coefficient that depends on the distance between the regions. Specifically, the (α, β) th element of Σ_2 is $\rho^{\text{dist}(q_\alpha, q_\beta)}$ for $\alpha, \beta \in \mathcal{P}$. Figure 11 illustrates the different correlation settings of Σ_1 and Σ_2 for region 25, which correspond to the 25th rows of Σ_1 and Σ_2 , respectively. We




Figure 10: Outbreak clusters when $c_{\text{out}} = 25$ (outbreak regions are in gray)

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	$\rho/2$	ρ	$\rho/2$	0	0
0	0	ρ	1	ρ	0	0
0	0	$\rho/2$	ρ	$\rho/2$	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

$\rho^{\sqrt{18}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{10}}$	ρ^3	$\rho^{\sqrt{10}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{18}}$
$\rho^{\sqrt{13}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{5}}$	ρ^2	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{13}}$
$\rho^{\sqrt{10}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{2}}$	ρ	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{10}}$
ρ^3	ρ^2	ρ	1	ρ	ρ^2	ρ^3
$\rho^{\sqrt{10}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{2}}$	ρ	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{10}}$
$\rho^{\sqrt{13}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{5}}$	ρ^2	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{13}}$
$\rho^{\sqrt{18}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{10}}$	ρ^3	$\rho^{\sqrt{10}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{18}}$

Figure 11: Correlation settings of Σ_1 and Σ_2 for region 25

run our examples with spatial correlation parameter $\rho = 0.0, 0.2, 0.4, 0.7$. However, the eigenvalues of Σ_1 approach zero or become negative when $\rho \geq 0.5$, which results in an ill-conditioned Σ_1 . Thus, we restrict $\rho < 0.5$ for Σ_1 .

Note that the number of regions in a cluster can be smaller near the edge of the study area. In order to keep the same number of regions in each cluster, we have added dummy regions. For example, we actually simulate 9×9 regions to monitor 7×7 regions with $r = 2$. Moreover, we conduct 100,000 replications of each experiment to calibrate empirical control limits, and 10,000 replications to obtain empirical ARLs (either ARL_0 or ARL_1). All analytical control limits and ARL_1 s are calculated by solving Equation (3.2.1) for $S_t^{25,r}$.

Hereafter, we mainly focus on the experimental results based on the case with $p = 7 \times 7$ and $ARL_0 = 100$, while other results with additional parameters (e.g., $p = 20 \times 20$ or the target $ARL_0 = 370$) are reported in Appendix B.

3.3.2 Range of the Scan Radius

The performance of the JMCUSUM-F chart depends heavily on the accuracy of Equation (3.2.1). For this reason, we determine the practical range of the scan radius r by studying how the accuracy of Equation (3.2.1) changes as r increases in the

JMCUSUM-F chart. Note that the JMCUSUM-F chart essentially runs p MCUSUM charts at the same time and stops when any of the MCUSUM charts signals an out-of-control alarm. We calculate analytical control limits from Equation (3.2.1) when the overall target ARL_0 is 100, 370, or 700. The case $ARL_0 = 700$ is additionally considered in this section to study the impact of the target ARL_0 . Then, we record the ratios of the empirical ARL_0 to the target ARL_0 for a single MCUSUM chart and the JMCUSUM-F chart, respectively, for $r = 0, 1, \sqrt{2}, 2,$ and $\sqrt{5}$ when $\delta = 1$ and $\rho = 0$.

From Section 3.1, we know that one MCUSUM statistic is based on the cumulative sum of $\ell_i^{c,r} = \boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_{c,r}/2)$. We see that the single MCUSUM chart is the same as a classical CUSUM chart with a reference value $k = 0.5\sqrt{\boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{c,r}}$ by setting $\ell_i^{c,r} = \boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \mathbf{x}_i - k\sqrt{\boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{c,r}}$, where $\sqrt{\boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{c,r}}$ is the standard deviation of $\ell_i^{c,r}$. For $\delta > 0$ and $0 \leq \rho < 1$, a large value of r increases the standard deviation of $\ell_i^{c,r}$; and since k is proportional to the standard deviation of $\ell_i^{c,r}$, k also increases as r increases. Table 9 illustrates the relationship between r and k when $\delta = 1$ and $\rho = 0$.

Figure 12 depicts the impact of r on the accuracy of Equation (3.2.1) for a single MCUSUM chart with various target ARL_0 s. It shows that the ratio of the empirical ARL_0 to the target ARL_0 decreases until r reaches 2. Figure 13 shows that the ratio decreases as the reference value k for a classical CUSUM chart increases up to about $k = 1.8$; Rogerson [44] also reports a similar decreasing trend as k increases for a

Table 9: r versus k when $\delta = 1$ and $\rho = 0$

r	k
0	0.5
1	≈ 1.1
$\sqrt{2}$	1.5
2	≈ 1.8
$\sqrt{5}$	≈ 2.3




Figure 12: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for a single MCUSUM chart




Figure 13: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus k for a single MCUSUM chart

classical CUSUM chart. One can easily see that Figures 12 and 13 match well. For example, the ratio of ARL₀s is approximately 80% when $r = \sqrt{2}$ in Figure 12, and the ratio is also about 80% when $k = 1.5$ in Figure 13.

The ratio graphs having different target ARL₀s start to behave quite differently when $r > 2$ and $k > 2$ as shown in Figures 12 and 13, respectively. From Figure 12, there can be more than a 30% drop in ARL₀ from the target when $r > 2$. However, to determine the range of r values where the approximation works well, we need to




Figure 14: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}(\%)$ versus r for the JMCUSUM-F chart

consider the performance of the JMCUSUM-F chart where dependence among the MCUSUM statistics exists.

In spatiotemporal monitoring, dependence among MCUSUM statistics is natural due to overlapping regions among clusters or spatial correlation. The dependence among the MCUSUM statistics becomes stronger as r or ρ increase. Figure 14 depicts the performance of the JMCUSUM-F chart with $p = 7 \times 7$, $\delta = 1$, $\rho = 0$, and target $\text{ARL}_0 = 100, 370$, and 700 . By comparing Figures 12 and 14, we see that the positive dependence among the MCUSUM statistics (due to overlapping regions) tends to shift the graphs of ARL_0 ratios in Figure 12 in the upward direction while keeping their overall shapes over r . We also notice that the tendency strengthens as r increases.

Figure 14 shows that the ratios are within a 10% difference of the target for various ARL_0 s when r is small (i.e., $0, 1, \sqrt{2}$), but the ratios deviate more from the target when r is large. This implies that the target ARL_0 does not significantly affect the accuracy of the approximation method if k is small. By considering the upward trend due to dependence and the impact of the target ARL_0 , we recommend the use of r such that $k \leq 2$ when $p = 7 \times 7$ and $\rho = 0$. Now, we explain the impacts of ρ and p .

- Spatial correlation ρ changes the standard deviation of $\ell_i^{c;r}$ and thus the values

of k . Therefore, for a given $\rho > 0$, we recommend that one ought to calculate k and find the range of r such that $k \leq 2$. For example, if $\rho = 0.2$, then $k \approx 1.8$ for $r = \sqrt{8}$, and $k \approx 2$ for $r = 3$ under the correlation structure Σ_2 . The ratios of the empirical ARL_0 to the target are 90.62% and 84.66%, respectively. Thus, in this case, one can use r up to 3 expecting no more than a 20% drop in the empirical ARL_0 from the target.

- Appendix B provides figures similar to Figures 12 and 14 for the case $p = 20 \times 20$. The figures show that there is still a shift in the upward direction, but the shift amount is smaller for $p = 20 \times 20$ when compared to $p = 7 \times 7$. This is because there are more MCUSUM statistics with fewer overlapping regions, and overall dependence among MCUSUM statistics diminishes with large p . Thus, the range of $k \leq 2$ is still recommended.

Following the recommendation, we test r up to 2 (i.e., $r_u = 2$), which ensures $k \leq 2$ for all values of ρ considered; and we set the largest $r_{out} = 2$ as well. As a consequence, we test the SMCUSUM chart with $r \in \mathcal{R} = \{0, 1, \sqrt{2}, 2\}$.

3.3.3 Results for the Fixed Scan Radius Case

We summarize some of our main results.

Control limits: Empirically and analytically evaluated control limits are presented in Table 10 under different configurations with a fixed scan radius. Although the $S_t^{c,r}$ s for different spatial clusters are no longer independent if $r > 0$ or $\rho > 0$, the empirical control limits match the analytical control limits well. The analytical approximation exhibits excellent performance, especially when $\rho < 0.5$.

In-control average run length: Before delivering test results for ARL_1^{JF} s, we need to confirm that the targeted ARL_0 is truly achieved with the analytical control limits. Empirical ARL_0 s of the JMCUSUM-F chart with analytical h^{**} are reported in Table 11. No matter what correlation structure is used, the results in Table 11

Table 10: Empirical h^{**} (Analytical h^{**}) when r is fixed

ρ	Σ_1			Σ_2		
	$r = 0$	$r = 1$	$r = 2$	$r = 0$	$r = 1$	$r = 2$
0.0	6.55 (6.64)	6.79 (6.81)	6.28 (6.17)	6.56 (6.64)	6.79 (6.81)	6.28 (6.17)
0.2	6.63 (6.70)	6.86 (6.86)	6.62 (6.52)	6.62 (6.70)	6.86 (6.86)	6.66 (6.55)
0.4	6.83 (6.86)	6.78 (6.74)	6.58 (6.38)	6.77 (6.80)	6.85 (6.82)	6.70 (6.56)
0.7	–	–	–	6.89 (6.87)	6.65 (6.43)	6.29 (5.94)

Table 11: Empirical ARL_0 of the JMCUSUM-F chart with analytical h^{**}

ρ	Σ_1			Σ_2		
	$r = 0$	$r = 1$	$r = 2$	$r = 0$	$r = 1$	$r = 2$
0.0	107.82	103.29	91.55	107.80	99.56	90.74
0.2	106.85	101.34	90.98	106.53	99.51	90.42
0.4	104.05	93.49	83.41	104.72	97.96	86.01
0.7	–	–	–	98.33	80.28	72.13

provide values close to 100 (i.e., the targeted ARL_0) when $\rho < 0.5$. The case in which $\rho > 0.5$ is tested as well for Σ_2 . However, the approximation quality seems diminished for such high ρ compared to the results for $\rho < 0.5$. As ρ and r become larger, stronger positive correlation occurs between spatial clusters and therefore smaller ARL_0 s are observed.

Out-of-control average run length: In Section 3.2.1, we suggest an analytical method to approximate ARL_1 s for the fixed scan radius case. Analytical ARL_1^{JF} s are compared to empirical ARL_1^{JF} s in Table 12. To see the gap caused by the ARL_1 approximation but not by the control limit approximation, analytical control limits are used to obtain both empirical and analytical ARL_1^{JF} s. Although the JMCUSUM-F chart yields some deviations of ARL_0 s from 100 (Table 11), ARL_1^{JF} s for both methods exhibit no significant difference (Table 12).

Additional ARL results for the JMCUSUM-F chart with $p = 7 \times 7$ and 20×20 , target $ARL_0 = 100, 370$, and $\rho = 0.0, 0.2, 0.4, 0.7$ are provided in Appendix B.

Table 12: Empirical ARL_1^{JF} (Analytical ARL_1^{JF}) with analytical h^{**}

ρ	Σ_1			Σ_2		
	$r = r_{out} = 0$	$r = r_{out} = 1$	$r = r_{out} = 2$	$r = r_{out} = 0$	$r = r_{out} = 1$	$r = r_{out} = 2$
0.0	13.34 (13.62)	3.33 (3.37)	1.49 (1.44)	13.29 (13.61)	3.37 (3.36)	1.47 (1.44)
0.2	11.66 (11.86)	4.06 (4.04)	2.07 (2.05)	11.82 (12.05)	4.04 (4.02)	2.15 (2.11)
0.4	6.88 (6.96)	2.88 (2.86)	1.81 (1.74)	8.88 (9.03)	3.54 (3.52)	2.21 (2.15)
0.7	–	–	–	4.44 (4.37)	1.93 (1.84)	1.33 (1.22)

3.3.4 Results for the Variable Scan Radius Case

In Section 3.2.2, we presented the SMCUSUM chart using a separate analytical control limit h^{*r} instead of a common control limit h^{**} . In order to compare the performance of the SMCUSUM chart to that of the JMCUSUM-V chart, ARL_0 s and ARL_1 s for the SMCUSUM chart are reported in Tables 13 and 14.

In Table 13, the SMCUSUM chart results in actual ARL_0 s close to the target $ARL_0 = 100$. Similar to the fixed scan radius case, smaller ARL_0 s are observed in Table 13 as the spatial correlation parameter, ρ , increases. In Table 14, ARL_1^{JV} (ARL_1^S) represents an empirical ARL_1 achieved by the JMCUSUM-V (SMCUSUM) chart from 10,000 replications. Although the SMCUSUM chart gives slightly larger (smaller) ARL_0 s for $\rho < 0.5$ ($\rho > 0.5$) than the target ARL_0 (Table 13), its ARL_1 performance is very close to that of the JMCUSUM-V chart (Table 14). Thus, the SMCUSUM chart provides similar performance to the JMCUSUM-V chart, but allows for analytical control limits — so that it is more convenient to implement the SMCUSUM chart than the JMCUSUM-V chart.

Next, we confirm that ARL_1^S is similar to the minimum of the ARL_1 s associated with the JMCUSUM-F charts, each applied with $r = 0, 1, \sqrt{2}, 2$, respectively. Both empirical and analytical ARL_1^{JF} s of the JMCUSUM-F charts are obtained for outbreak scenarios with $r_{out} = 0, 1, 2$. For each value of r_{out} , we note that the JMCUSUM-F chart with $r = r_{out}$ gives the minimum of the ARL_1^{JF} s. The minimum ARL_1^{JF} s are given in Table 12; and the ARL_1^{JF} s in Table 12 are very close to the ARL_1^S s in Table

Table 13: Empirical ARL_0 for the SMCUSUM chart with analytical h^{*r}

ρ	Σ_1	Σ_2
0.0	117.66	118.82
0.2	113.68	113.11
0.4	105.77	106.40
0.7	–	87.13

Table 14: Empirical ARL_1^{JV} (ARL_1^S)

ρ	Σ_1			Σ_2		
	$r_{out} = 0$	$r_{out} = 1$	$r_{out} = 2$	$r_{out} = 0$	$r_{out} = 1$	$r_{out} = 2$
0.0	14.60 (15.05)	3.66 (3.79)	1.67 (1.73)	14.50 (15.23)	3.62 (3.76)	1.72 (1.73)
0.2	13.40 (13.23)	4.60 (4.62)	2.42 (2.60)	13.30 (13.56)	4.48 (4.61)	2.50 (2.67)
0.4	8.03 (7.88)	3.39 (3.33)	2.20 (2.47)	10.10 (10.21)	4.08 (4.12)	2.60 (2.89)
0.7	–	–	–	4.94 (5.01)	2.26 (2.22)	1.59 (1.81)

14. For example, when Σ_1 is used with $r_{out} = 1$ and $\rho = 0.2$, the minimum empirical (analytical) ARL_1^{JF} is 4.06 (4.04), which is only slightly smaller than the empirical ARL_1^{JV} (ARL_1^S), 4.60 (4.62). The ARL_1^S s tend to be slightly larger than the minimum ARL_1^{JF} s for the configurations we tested, but the differences are not significant, supporting our conjecture in Section 3.2.2. The same tendency is observed for ARL_1^{JV} s as well.

Additional ARL results for the SMCUSUM chart are provided in Appendix B for $p = 7 \times 7$ and 20×20 , with target $ARL_0 = 100, 370$ and $\rho = 0.0, 0.2, 0.4, 0.7$.

3.4. Impact of Spatial Correlation on Detection Performance

This section investigates the impact of correlation on detection performance as measured by detection delay and spatial identification accuracy, using the JMCUSUM-F and SMCUSUM charts.

For a control chart used in spatiotemporal biosurveillance, two important questions are: (i) how quickly does the chart detect an outbreak? and (ii) how accurate is

the chart in identifying the outbreak cluster? ARL_1 s are a generally accepted measure of detection delay and thus can be used to answer the first question.

To answer the second question, we focus on two spatial identification accuracy measures: the correct alarm percentage (CP) giving the proportion of time that a chart exactly detects the true outbreak cluster with the correct center and radius; and the sum of the alarm percentages (SP) that detect a cluster including all of the outbreak regions — in the context of contagious disease surveillance, the outcome we most want to avoid is that of missing any outbreak region. Other identification accuracy measures for spatial cluster detection can be found in the literature. For example, Gangnon and Clayton [15] report the proportion of time that their methods detect at least one, half of, or all outbreak regions in the outbreak cluster. Gangnon and Clayton [16] also discuss the power of their likelihood ratio tests (i.e., cluster detection rates from known clustering models).

In order to investigate the relationship between detection performance and spatial correlation, we run a series of experiments under various scenarios. The detection performance is tested under the same experimental setup from Section 3.3.1, but with one additional value of $\delta = 0.5$. We use the JMCUSUM-F chart from Section 3.2.1 for the fixed scan radius case, and the SMCUSUM chart from Section 3.2.2 for the variable scan radius case. As a measure of detection delay, ARL_1 s for each configuration are presented in the last four columns of Tables 15 and 16. All of the ARL_1 s reported in the tables are empirically obtained by simulation with analytical control limits based on 10,000 replications. As measures of identification accuracy, Tables 15 and 16 report values of CP and SP under the correlation structures Σ_1 and Σ_2 , respectively.

Table 15: Identification accuracy CP (SP) with Σ_1

δ	ρ	r_{out}	CP (SP)				Empirical ARL ₁				
			$r = 0$	$r = 1$	$r = 2$	$r = \text{var}$	$r = 0$	$r = 1$	$r = 2$	$r = \text{var}$	
0.0	0.0	0	83.73 (83.73)	– (49.61)	– (47.47)	63.93 (81.63)	40.02	71.34	93.75	44.87	
		1	–	85.89 (85.89)	– (73.06)	68.74 (83.59)	22.85	10.67	18.75	12.27	
		2	–	–	73.89 (73.89)	65.52 (65.52)	17.94	6.90	4.63	5.67	
	0.5	0.2	0	86.81 (86.81)	– (52.19)	– (50.84)	66.91 (83.69)	35.53	70.67	88.82	40.04
		1	–	89.16 (89.16)	– (69.49)	77.30 (85.86)	28.15	12.94	25.81	15.09	
		2	–	–	83.09 (83.09)	76.17 (76.17)	24.29	11.38	6.62	8.35	
1.0	0.4	0	91.26 (91.26)	– (62.88)	– (57.73)	74.17 (89.39)	21.90	52.27	73.64	25.65	
		1	–	89.65 (89.65)	– (65.15)	85.73 (89.95)	18.64	9.26	19.10	10.96	
		2	–	–	87.26 (87.26)	78.93 (78.93)	17.48	10.21	5.76	7.77	
	0.0	0	95.40 (95.40)	– (61.65)	– (51.16)	82.24 (91.97)	13.34	46.18	59.90	15.05	
		1	–	92.39 (92.39)	– (80.73)	79.05 (88.62)	7.65	3.33	7.23	3.79	
		2	–	–	81.89 (81.89)	71.42 (71.42)	6.10	2.26	1.47	1.73	
1.0	0.2	0	96.14 (96.14)	– (64.42)	– (56.46)	85.40 (93.71)	11.66	45.17	62.48	13.23	
		1	–	95.43 (95.43)	– (78.38)	87.72 (92.84)	9.68	4.06	10.55	4.62	
		2	–	–	89.59 (89.59)	81.18 (81.18)	8.34	3.74	2.07	2.60	
	0.4	0	96.99 (96.99)	– (76.44)	– (62.57)	87.76 (96.04)	6.88	27.57	45.76	7.88	
		1	–	94.66 (94.66)	– (73.62)	92.52 (94.74)	6.14	2.88	7.41	3.33	
		2	–	–	91.50 (91.50)	81.77 (81.77)	5.97	3.36	1.81	2.47	

Table 16: Identification accuracy CP (SP) with Σ_2

δ	ρ	r_{out}	CP (SP)				Empirical ARL ₁				
			$r = 0$	$r = 1$	$r = 2$	$r = \text{var}$	$r = 0$	$r = 1$	$r = 2$	$r = \text{var}$	
0.0	0.0	0	84.17 (84.17)	– (49.56)	– (47.64)	63.71 (81.02)	40.40	72.06	94.72	44.87	
		1	–	85.93 (85.93)	– (72.41)	68.11 (82.69)	22.70	10.70	18.68	12.40	
		2	–	–	73.48 (73.48)	64.55 (64.55)	17.88	6.90	4.65	5.65	
	0.5	0.2	0	86.88 (86.88)	– (52.41)	– (53.64)	67.43 (83.69)	36.16	71.24	89.39	41.08
		1	–	90.60 (90.60)	– (71.55)	79.75 (87.79)	28.16	12.91	26.13	15.15	
		2	–	–	85.67 (85.67)	78.98 (78.98)	24.79	12.04	6.82	8.57	
1.0	0.4	0	89.90 (89.90)	– (56.53)	– (57.98)	72.01 (87.40)	27.80	64.37	78.95	31.82	
		1	–	92.69 (92.69)	– (69.53)	87.19 (91.95)	25.46	11.50	28.49	13.52	
		2	–	–	91.48 (91.48)	84.84 (84.84)	22.69	15.07	7.12	9.22	
	0.0	0	94.75 (94.75)	– (60.90)	– (52.04)	82.66 (91.90)	13.29	46.06	62.09	15.23	
		1	–	92.70 (92.70)	– (81.68)	79.06 (88.65)	7.57	3.37	7.24	3.76	
		2	–	–	82.00 (82.00)	70.93 (70.93)	6.03	2.25	1.47	1.73	
1.0	0.2	0	95.62 (95.62)	– (65.51)	– (56.40)	85.17 (93.69)	11.82	45.85	63.03	13.56	
		1	–	96.08 (96.08)	– (81.14)	88.50 (93.57)	9.62	4.04	11.24	4.61	
		2	–	–	91.70 (91.70)	83.22 (83.22)	8.55	4.01	2.15	2.67	
	0.4	0	97.18 (97.18)	– (70.15)	– (63.50)	88.22 (95.46)	8.88	38.39	54.55	10.21	
		1	–	97.09 (97.09)	– (79.22)	93.61 (96.07)	8.74	3.54	12.44	4.12	
		2	–	–	95.31 (95.31)	87.27 (87.27)	7.82	5.13	2.21	2.89	

3.4.1 Outbreak Radius

With respect to the ARL_1 s in Tables 15 and 16, it is easy to see the importance of matching the scan radius to the actual outbreak radius. The use of the fixed scan radius $r = r_{out}$ always results in the smallest ARL_1 s and the largest alarm percentages under all configurations tested. The SMCUSUM chart with a variable scan radius always provides the second-best ARL_1 s. With respect to spatial identification accuracy, the alarm percentages when using the variable scan radius are not as good as those with the fixed scan radius $r = r_{out}$, but give reasonably high values, especially when the correlation is high. If the radius of the real outbreak increases, then both the JMCUSUM-F and SMCUSUM charts raise an alarm more quickly. However, the values of CP and SP decrease as r_{out} increases. There is not much difference between the experimental results for Σ_1 and Σ_2 , except that the ARL_1 s decrease more slowly as ρ increases when Σ_2 is used.

3.4.2 Spatial Correlation

To see the impact of spatial correlation on both the ARL_1 and alarm percentages more clearly, we depict the results of the fixed scan radius case. Figure 15 shows the ARL_1 and CP as functions of correlation for $r_{out} = 0, 1, 2$ with $\delta = 1$ and Σ_2 . In Figure 15, the ARL_1 does not have a monotonic relationship with spatial correlation, but the CP increases as correlation increases. This shows that positive correlation does not always support faster detection but helps accurate identification. (The same ARL_1 and CP patterns appear with the variable scan radius.) Regarding the relationship between spatial correlation and ARL_1 s, one can find a similar conclusion in Rogerson and Yamada [46], although the authors use different MCUSUM charts than ours. We drew graphs of ARL_1 and SP versus spatial correlation with different δ and Σ_1 values, which are not reported in this thesis, and we observed similar patterns. These results imply that faster detection does not guarantee accurate identification.




Figure 15: Spatial correlation vs. ARL_1 and CP with $\delta = 1$ and Σ_2

3.5. Conclusions

We propose a method to calculate approximate control limits for the JMCUSUM-F chart where the scan radius is fixed, and to study the range of the scan radius where the method works well. A new MCUSUM chart is also proposed for the case of a variable scan radius, namely, the SMCUSUM chart, whose performance is similar to the JMCUSUM-V chart but calculates control limits analytically. By using the proposed charts, one can avoid cumbersome time-consuming trial-and-error simulation to calibrate control limits, which makes the use of MCUSUM charts more convenient in practice.

In addition, we introduce the correct alarm percentage and the sum of alarm percentages to evaluate identification accuracy of the proposed charts for biosurveillance; and we study the impact of the outbreak radius and spatial correlation on ARL_1 s and identification accuracy.

The experimental results show that matching r to r_{out} is important not only with regard to ARL_1 s, but also for identification accuracy. Using a chart designed for the variable scan radius case is a safer choice for both ARL_1 and identification accuracy

if r_{out} is unknown. Larger r_{out} helps fast detection but makes identification of the outbreak cluster difficult. On the other hand, we find that higher spatial correlation does not always help faster detection, but it does help accurate identification.

Chapter IV

ROBUST DISTRIBUTION-FREE MCUSUM CHARTS FOR SPATIOTEMPORAL BIOSURVEILLANCE IN THE PRESENCE OF SPATIAL CORRELATION

In Chapter 3, we focused on an analytical method to approximate the control limits in order to overcome the fact that many MCUSUM charts rely on time-consuming trial-and-error simulations to search for their control limits. Based on the MCUSUM chart developed by Jiang et al. [24] (the JMCUSUM chart), we suggested another MCUSUM chart, named the SMCUSUM chart, that applies analytical control limits separately to each of the group of MCUSUM charts having the same scan radius. We tested our method under various configurations of outbreak coverage, shift magnitude, spatial correlations among regions, and scan radius. The control limit approximation method works well in a certain range of radius values, while the ARL_1 performance of the SMCUSUM charts is similar to that of the JMCUSUM charts.

In this chapter, we design and develop a new MCUSUM chart, namely, the Robust-MCUSUM (RMCUSUM) chart, motivated from the JMCUSUM and SMCUSUM charts. Our new method searches for the control limits of the correlated MCUSUM charts for any underlying distribution of data (i.e., distribution-free) and any range of scan radius. The chapter is organized as follows: Section 4.1 defines notation and the problem itself, reviews the monitoring statistics of the JMCUSUM and SMCUSUM charts, and gives motivating examples. Section 4.2 presents our new RMCUSUM chart, and then Section 4.3 delivers the experimental setup and results to test the performance of the RMCUSUM chart. Section 4.4 concludes this chapter.

4.1. Background and Motivation

As background, we first define the notation needed to describe the problem and the RMCUSUM chart. Then we explain how our motivation arises from the JMCUSUM and SMCUSUM charts.

4.1.1 Notation and Problem

We consider $p = M \times N$ regions in a rectangular shape for monitoring. The two-dimensional space coordinate of region c in a set $\mathcal{P} = \{1, 2, \dots, p\}$ is denoted as $q_c = (m_c, n_c)$, where $m_c = 1, \dots, M$ and $n_c = 1, \dots, N$ (See Figure 16). At each time t , an observation vector $\mathbf{y}_t = (y_{t1}, y_{t2}, \dots, y_{tp})'$ is monitored from p regions. If the baseline mean value in the in-control state, θ_c , and the marginal variance, σ_c^2 , of each component y_{tc} are known for all $c \in \mathcal{P}$, we obtain a standardized observation vector $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$ by setting $x_{tc} = (y_{tc} - \theta_c)/\sigma_c$. The quantity \mathbf{x}_t is assumed to follow a multivariate distribution having probability density function $f_{\boldsymbol{\mu}}(\mathbf{x})$ with a mean vector $\boldsymbol{\mu}$ and known variance-covariance matrix Σ . We assume that Σ does not change over times $t = 1, 2, \dots$, while $\boldsymbol{\mu}$ can change when there is a shift. With the standardized observation vectors, the variance-covariance matrix and the correlation matrix are interchangeable.

Since each component of \mathbf{x}_t corresponds to a standardized observation from each spatial region at time t , we can construct an MCUSUM chart on \mathbf{x}_t (or a subset of \mathbf{x}_t) for spatiotemporal biosurveillance. The MCUSUM chart is developed to detect a shift as soon as possible after the null parameter $\boldsymbol{\mu}_0 = \mathbf{0}$ in the in-control state is shifted to the alternative parameter $\boldsymbol{\mu}_1$ (which indicates an out-of-control state) at an unknown change time ν .

Employing the assumption from Jiang et al. [24] that the shape of the outbreak coverage is a circle, we define possible spatial clusters of outbreak regions as $O^{c,r} \equiv \{j \mid \text{dist}(q_j, q_c) \leq r, j \in \mathcal{P}\}$ with scan radius r from region c , where $\text{dist}(a, b)$ denotes

c						
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	32	33	34	35
36	37	38	39	40	41	42
43	44	45	46	47	48	49

$q_c = (m_c, n_c)$						
(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)	(1, 7)
(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)	(2, 7)
(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)	(3, 7)
(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)	(4, 7)
(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)	(5, 7)
(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)	(6, 7)
(7, 1)	(7, 2)	(7, 3)	(7, 4)	(7, 5)	(7, 6)	(7, 7)

Figure 16: Coordinate expression of regions ($p = 7 \times 7$)

the Euclidean distance between a and b . Since the center of the circle-shaped coverage is located in c , let c denote the center region of the outbreak cluster $O^{c,r}$.

The scan radius r can be either constant or variable. When r is variable, r can be in a range from $[0, \sqrt{M^2 + N^2}]$ or bounded by a given upper limit r_u . However, since the number of regions in $O^{c,r}$ changes only at certain values of r , we consider a finite number of possible values of r such as 0, 1, $\sqrt{2}$, 2, and so on for practical purposes. Let \mathcal{R} denote the set of possible settings for r , i.e., $\mathcal{R} = \{r_1, r_2, \dots, r_u\}$, where u is the cardinality of \mathcal{R} .

We scan each possible outbreak cluster $O^{c,r}$ and test if the mean levels of the regions in $O^{c,r}$ are shifted, assuming the mean levels of the other regions stay in the in-control state. For different $O^{c,r}$ choices, different shift vectors $\boldsymbol{\mu}_{c,r}$ are considered, and a set of possible shift vectors is defined as $\{\boldsymbol{\mu}_{c,r} \mid c \in \mathcal{P}, r \in \mathcal{R}\}$. When a homogeneous shift magnitude $\delta > 0$ is assumed over all outbreak regions, the j th component of $\boldsymbol{\mu}_{c,r}$ is δ if $j \in O^{c,r}$ and 0 otherwise. An MCUSUM chart is designed to detect a shift from $\boldsymbol{\mu}_0$ to $\boldsymbol{\mu}_{c,r}$ as soon as possible.

In advance of presenting the RMCUSUM chart in Section 4.2, we introduce some additional notation and assumptions.

Suppose the discrete-time stochastic process $\{Y_i : i = 1, 2, \dots\}$ has a steady-state distribution with marginal mean $E[Y_i]$ and marginal variance $\text{Var}[Y_i]$. Then the standardized CUSUM is defined as

$$\mathcal{C}_t(s) \equiv \frac{\sum_{i=1}^{\lfloor ts \rfloor} Y_i - tsE[Y_i]}{\Omega_Y \sqrt{t}} \quad \text{for } s \in [0, 1], \quad (4.1.1)$$

where $\lfloor \cdot \rfloor$ is the floor function and Ω_Y^2 is the variance parameter for the process $\{Y_i\}$, defined as

$$\Omega_Y^2 \equiv \lim_{t \rightarrow \infty} t \text{Var}[\bar{Y}(t)] = \sum_{h=-\infty}^{\infty} \text{Cov}(Y_i, Y_{i+h})$$

when $\bar{Y}(t) \equiv t^{-1} \sum_{i=1}^t Y_i$ is the sample mean of the first t observations. Note that $\Omega_Y^2 = \text{Var}[Y_i]$ if the process $\{Y_i : i = 1, 2, \dots\}$ has no autocorrelation over time.

Let $\{\mathcal{W}(u) : u \in [0, \infty)\}$ denote a standard Brownian motion process so that the random variables $\mathcal{W}(s)$ and $\mathcal{W}(u)$ for arbitrary $s, u \in [0, \infty)$ are jointly normal with $E[\mathcal{W}(s)] = E[\mathcal{W}(u)] = 0$ and $\text{Cov}[\mathcal{W}(s), \mathcal{W}(u)] = \min\{s, u\}$.

For each positive integer t , the random function $\mathcal{C}_t(\cdot)$ is an element of $D[0, 1]$, the space of functions on $[0, 1]$ that are right-continuous and have left-hand limits [4]. As Kim et al. [25] do, we assume that $\{Y_i : i = 1, 2, \dots\}$ satisfies the following Functional Central Limit Theorem (FCLT).

Assumption FCLT. *There exist finite real constants $E[Y_i]$ and $\Omega_Y^2 > 0$ such that the sequence of random functions $\{\mathcal{C}_t(\cdot) : t = 1, 2, \dots\}$ converges in distribution to a standard Brownian motion $\mathcal{W}(\cdot)$ in the space $D[0, 1]$ as $t \rightarrow \infty$. Formally, we write*

$$\mathcal{C}_t(\cdot) \xrightarrow[t \rightarrow \infty]{D} \mathcal{W}(\cdot),$$

where $\xrightarrow[t \rightarrow \infty]{D}$ denotes convergence in distribution as $t \rightarrow \infty$.

Note that, based on the definition (4.1.1) and the assumption FCLT, the following

approximation holds for Y_i in any distribution if t is sufficiently large.

$$\begin{aligned}
\sum_{i=1}^t Y_i &= \text{E}[Y_i]t + \Omega_Y \sqrt{t} \mathcal{C}_t(1) \\
&\stackrel{\text{D}}{\approx} \text{E}[Y_i]t + \Omega_Y \sqrt{t} \mathcal{W}(1) \\
&\stackrel{\text{D}}{=} \text{E}[Y_i]t + \Omega_Y \mathcal{W}(t)
\end{aligned} \tag{4.1.2}$$

where $\stackrel{\text{D}}{\approx}$ denotes approximate equality in distribution and $\stackrel{\text{D}}{=}$ denotes exact equality in distribution.

As we saw in Chapter 3, both of the JMCUSUM and SMCUSUM charts are based on the monitoring statistic,

$$S_t^{**} \equiv \max_{c,r} S_t^{c,r} \equiv \max_{c,r} \max_{1 \leq \nu^* \leq t} \sum_{i=\nu^*}^t \ell_i^{c,r}, \quad t = 1, 2, \dots,$$

where

$$\ell_i^{c,r} \equiv \boldsymbol{\mu}'_{c,r} \Sigma^{-1} \left(\mathbf{x}_i - \frac{\boldsymbol{\mu}_{c,r}}{2} \right), \quad i = 1, 2, \dots, t.$$

While the JMCUSUM chart applies a control limit directly to the monitoring statistic, the SMCUSUM chart applies separate control limits according to different scan radius options.

Note that $S_t^{c,r}$ itself is an MCUSUM statistic that can be calculated recursively by

$$S_t^{c,r} = \max(0, S_{t-1}^{c,r} + \ell_t^{c,r}), \quad t = 1, 2, \dots,$$

where $S_0^{c,r} = 0$ for all $c \in \mathcal{P}$ and $r \in \mathcal{R}$. Then it is easy to see that $S_t^{c,r}$ is always equal to $S^{c,r}(t) - \min\{S^{c,r}(n) : n = 1, 2, \dots, t\}$, where $S^{c,r}(t) \equiv \sum_{i=1}^t \ell_i^{c,r}$.

4.1.2 Motivation

Though both of the JMCUSUM and SMCUSUM charts perform well, the SMCUSUM chart has an advantage. While searching for the appropriate control limit that yields the target ARL_0 , the JMCUSUM chart uses tedious trial-and-error simulation that requires intensive modeling of the underlying process. The SMCUSUM chart uses an

analytical approximation method that is much faster in terms of calculating its control limit than trial-and-error simulation. In spite of the advantage, we still present some motivations for developing a new MCUSUM chart for practical use.

First, Chapter 3 showed that the performance of the SMCUSUM chart highly depends on a reference value for the chart (i.e., $k = 0.5\sqrt{\boldsymbol{\mu}'_{c,r}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}_{c,r}}$) and recommends use of a reference value $k \leq 2$. With this condition on the reference value, the range of the scan radius is restricted, and this restricted range may cause inconvenience for users, especially when the monitoring area is large. In addition, since the performance of the SMCUSUM chart deteriorates as the reference value increases, that control limit approximation method becomes less reliable when the reference value gets closer to 2. For example, when the number of regions $p = 7 \times 7$, the shift magnitude $\delta = 1$, the scan radius $r = 2$, and no spatial correlation exists among the regions, the reference value of the SMCUSUM chart is approximately 1.8. With the reference value close to 2, the empirical ARL_0 of the SMCUSUM chart stays around 80% of the target ARL_0 , 370.

Second, the control limit approximation method for the SMCUSUM chart is based on the independence assumption on the underlying MCUSUM statistics, $S_t^{c,r}$. However, the assumption is not valid if there is positive spatial correlation among the data or if there are overlapped regions among the spatial clusters.

Finally, both of the JMCUSUM and SMCUSUM charts are developed and tested under the assumption that the regional disease counts data follow a multivariate normal distribution. Unfortunately, the underlying distribution of incident data is often unknown or modeled by a Poisson distribution, and the performance of the charts originally developed for normal data may deteriorate when applied to non-normal data. For example, when we apply the JMCUSUM chart to multivariate Poisson data having all marginal means equal to 5 with the number of regions $p = 7 \times 7$, the shift magnitude $\delta = 1$, the scan radius $r = 0$, and the spatial correlation

parameter $\rho = 0.4$, the empirical ARL_0 drops to 44.61, which is far from the target $\text{ARL}_0 = 100$.

In Section 4.2, we introduce our RMCUSUM chart with a control limit search method that can handle the correlations among the MCUSUM charts without intensive modeling of the underlying process, while achieving robustness to any scan radius range, spatial correlation, and data distribution.

4.2. RMCUSUM Charts

For the RMCUSUM chart, we introduce the following monitoring statistic,

$$G_t^{**} \equiv \max_{c,r} G_t^{c,r} \equiv \max_{c,r} \max_{1 \leq \nu^* \leq t} \sum_{\ell=\nu^*}^t (\alpha_\ell^{c,r} - k\sigma^{c,r}), \quad t = 1, 2, \dots$$

where $\alpha_\ell^{c,r} \equiv \boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \mathbf{x}_\ell$, k is a constant reference value, and $\sigma^{c,r}$ is the standard deviation of $\alpha_\ell^{c,r}$ (i.e., $\sqrt{\boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{c,r}}$) for $\ell = 1, 2, \dots, t$. An alarm is signaled as soon as $G_t^{**} \geq g^{**}$, where g^{**} is an approximated control limit that determines the operating characteristics of the monitoring chart. When a signal is raised for the first time, we record (i) the current time t as the time when the control chart detects the shift and (ii) ν^* as the estimated change time, though the true change time ν is still unknown.

Similar to $S_t^{c,r}$, the quantity $G_t^{c,r}$ is itself an MCUSUM statistic that can be calculated by

$$\begin{aligned} G_t^{c,r} &= \max(0, G_{t-1}^{c,r} + \alpha_t^{c,r} - k\sigma^{c,r}) \\ &= S_c(t) - \min\{S_c(i) : i = 1, 2, \dots, t\}, \quad t = 1, 2, \dots, \end{aligned} \quad (4.2.1)$$

where $S_c(t) \equiv \sum_{\ell=1}^t (\alpha_\ell^{c,r} - k\sigma^{c,r})$ and $G_0^{c,r} = 0$ for all $c \in \mathcal{P}$ and $r \in \mathcal{R}$. Therefore, one MCUSUM statistic $G_t^{c,r}$ is recorded for each spatial cluster, and the maximum of all pu of the MCUSUM statistics is used as the monitoring statistic G_t^{**} .

It is notable that our monitoring statistic G_t^{**} becomes S_t^{**} when $k = 0.5\sqrt{\boldsymbol{\mu}'_{c,r} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{c,r}}$. Chapter 3 suggested an analytical control limit approximation

method for S_t^{**} that can be applicable while $k \leq 2$. Though that approximation method is very fast and convenient for calculating the control limit, the possible range of the scan radius r is restricted due to the constraint on k . In order to avoid such a restriction on r , we use a constant k that does not depend on r as in a classical CUSUM chart.

For the RMCUSUM chart, it is desirable that (i) g^{**} should deliver the exact target ARL_0 ; and (ii) the search for such g^{**} should not require cumbersome modeling of the underlying process, which can be challenging with large r . In the next two subsections, we present our RMCUSUM chart, which incorporates a new control limit search method that achieves these two goals when the scan radius r is fixed or variable.

4.2.1 RMCUSUM Chart for the Fixed Scan Radius

We first consider the RMCUSUM chart for the fixed scan radius, namely, RMCUSUM-F. If the actual outbreak radius r_{out} is known or can be approximated well, r can be fixed to r_{out} for good ARL_1 performance of the chart. When r is fixed, our monitoring statistic becomes

$$G_t^{**} = \max_c G_t^{c,r}, \quad t = 1, 2, \dots$$

This means that the number of possible spatial clusters $O^{c,r}$ is reduced to p , and one MCUSUM statistic, $G_t^{c,r}$, is monitored for each spatial cluster. An alarm is raised when any of the p MCUSUM statistics exceed a common control limit g^{**} . If an alarm is incurred by $G_t^{c,r}$, then c is considered to be the center of the outbreak.

Due to the overlapped regions among the $O^{c,r}$ s for $c \in \mathcal{P}$ and the spatial correlation existing in the regional data \mathbf{x}_t , it is typically true that the $G_t^{c,r}$ s for $c \in \mathcal{P}$ are correlated and that the correlation among the $G_t^{c,r}$ s makes the search for proper control limits difficult — as one can observe from the performance deterioration of the SMCUSUM chart in Chapter 3. We propose a method that deals with the correlation

among the $G_t^{c,r}$ s.

First, note that

$$\begin{bmatrix} G_t^{1,r} \\ G_t^{2,r} \\ \vdots \\ G_t^{p,r} \end{bmatrix} \equiv \begin{bmatrix} S_1(t) - \min\{S_1(i) : i = 1, 2, \dots, t\} \\ S_2(t) - \min\{S_2(i) : i = 1, 2, \dots, t\} \\ \vdots \\ S_p(t) - \min\{S_p(i) : i = 1, 2, \dots, t\} \end{bmatrix}$$

We define $\boldsymbol{\alpha}(\ell)$ as the $p \times 1$ vector $(\alpha_\ell^{1,r}, \dots, \alpha_\ell^{p,r})'$ and Γ as the $p \times p$ variance-covariance matrix of $\boldsymbol{\alpha}(\ell)$. Since $\alpha_\ell^{c,r}$ s for all $c \in \mathcal{P}$ are linear combinations of the data \boldsymbol{x}_ℓ , Γ can be derived analytically from the known Σ . When Γ is a positive-definite matrix, there exists $\Gamma^{\frac{1}{2}}$ that satisfies $\Gamma = \Gamma^{\frac{1}{2}}[\Gamma^{\frac{1}{2}}]'$. Then, for the data with any distribution,

$$\boldsymbol{S}(t) \equiv \begin{bmatrix} S_1(t) \\ S_2(t) \\ \vdots \\ S_p(t) \end{bmatrix} \stackrel{D}{\approx} -\mathbf{K}t + \Gamma^{\frac{1}{2}}\mathbf{W}(t), \quad (4.2.2)$$

where \mathbf{K} denotes a $p \times 1$ vector $(k\sigma^{1,r}, \dots, k\sigma^{p,r})'$, $\mathbf{W}(t)$ denotes a $p \times 1$ vector $(\mathcal{W}_1(t), \dots, \mathcal{W}_p(t))'$ with independent standard Brownian motion process components (see Appendix C for the detailed derivation). Note that Equation (4.2.2) is the multivariate version of Equation (4.1.2). By Equation (4.1.2), the values of $\mathcal{W}_c(t)$ for $c \in \mathcal{P}$ and integer values of t can be generated by the summation of t i.i.d. standard normal observations.

We search by trial-and-error for g^{**} yielding the target ARL_0 . Figure 17 provides the detailed procedure of the RMCUSUM-F chart, including the trial-and-error simulation search for its control limit. In Figure 17, Phase I depicts the steps for finding the control limit value g^{**} , while Phase II shows the step for running the RMCUSUM-F chart with the resulting g^{**} . Though g^{**} is determined by trial-and-error, the search does not require any complicated modeling to generate correlated

Phase I. *Searching for the control limit.*

- Step 1. Select a value of g^{**} and determine the number of replications τ .
- Step 2. Set $t = 1$, $\mathbf{S}(0) = \mathbf{0}$, $\mathbf{W}(0) = \mathbf{0}$, and the iteration counter $i = 0$.
- Step 3. Generate $\mathbf{Z} = [Z_1, Z_2, \dots, Z_p]'$ where the Z_c s are i.i.d. standard normal random variables for all $c \in \mathcal{P}$.
- Step 4. Set $\mathbf{W}(t) = \mathbf{W}(t-1) + \mathbf{Z}$ and $\mathbf{S}(t) = -\mathbf{K}t + \Gamma^{\frac{1}{2}}\mathbf{W}(t)$.
- Step 5. For $c \in \mathcal{P}$, set

$$\begin{aligned} G_t^{c,r} &= S_c(t) - \min\{S_c(i) : i = 1, \dots, t\} \\ &= \max\left(0, G_{t-1}^{c,r} - k\sigma^{c,r} + [\Gamma^{\frac{1}{2}}\mathbf{Z}]_c\right) \end{aligned}$$

where $[\Gamma^{\frac{1}{2}}\mathbf{Z}]_c$ denotes the c th element of $\Gamma^{\frac{1}{2}}\mathbf{Z}$.

- Step 6. If $G_t^{**} \geq g^{**}$, raise an alarm, record t as a run length, and set $i = i + 1$.
- Step 7. If $i = \tau$, go to Step 8. Otherwise, set $t = t + 1$ and go to Step 3.
- Step 8. Calculate ARL_0 . If it is close to the target ARL_0 , go to Phase II. Otherwise, pick a different value for g^{**} and go to Step 2.

Phase II. *Running the RMCUSUM_f chart with the control limit.*

- Step 1. For each $t = 1, 2, \dots$, obtain \mathbf{x}_t and calculate G_t^{**} . Raise an out-of-control alarm if $G_t^{**} \geq g^{**}$.

Remark: To select a value of g^{**} efficiently in Phase I, one may use a bisection method or other metaheuristics instead of trying arbitrary values.

Figure 17: The RMCUSUM_f chart

variables — we only need to generate p i.i.d. standard normal variables.

4.2.2 RMCUSUM Chart for the Variable Scan Radius

If r_{out} is unknown, the use of a variable r is recommended. In this case, G_t^{**} becomes equivalent to the maximum of pu statistics, and an alarm is raised as soon as any of the pu MCUSUM statistics hits a common control limit g^{**} .

In order to find g^{**} for the pu correlated MCUSUM charts, we can consider a chart similar to the RMCUSUM-F by expanding the dimension from p to pu . However, with

the expanded dimension pu , Γ is often not positive-definite, which implies $\Gamma^{\frac{1}{2}}$ does not exist. Therefore, instead of using the simple expansion of the RMCUSUM-F chart, we introduce an alternative chart RMCUSUM-V, which is similar to the SMCUSUM chart.

Before discussing the RMCUSUM-V chart, we need to define two types of correlation related to the overlapped regions of spatial clusters for the pu MCUSUM charts. One is the correlation among the charts with the same radius (SR-type correlation). The other is the correlation among the charts with different radii (DR-type correlation). For example, G_t^{1,r_2} and G_t^{2,r_2} have SR-type correlation, while G_t^{1,r_2} and G_t^{2,r_3} have DR-type correlation.

Since the correlation among the $\alpha_t^{c,r}$ s directly affects the correlation among the $G_t^{c,r}$ s, the correlation among the pu MCUSUM charts can be presented by a $pu \times pu$ matrix

$$\Psi \equiv \begin{bmatrix} \Psi_{r_1,r_1} & \Psi_{r_1,r_2} & \cdots & \Psi_{r_1,r_{u-1}} & \Psi_{r_1,r_u} \\ \Psi_{r_2,r_1} & \Psi_{r_2,r_2} & \cdots & \Psi_{r_2,r_{u-1}} & \Psi_{r_2,r_u} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \Psi_{r_u,r_1} & \Psi_{r_u,r_2} & \cdots & \Psi_{r_u,r_{u-1}} & \Psi_{r_u,r_u} \end{bmatrix}$$

where the (i,j) th element of a $p \times p$ matrix Ψ_{r_a,r_b} is the correlation between α_t^{i,r_a} and α_t^{j,r_b} for $i, j \in \mathcal{P}$ and $r_a, r_b \in \mathcal{R}$. Note that Ψ_{r_a,r_b} with $r_a = r_b$ is related to the SR-type correlation and Ψ_{r_a,r_b} with $r_a \neq r_b$ is related to the DR-type correlation.

Recall that the SMCUSUM chart applies different control limits to the MCUSUM charts with different r (while the JMCUSUM chart applies one common control limit to all MCUSUM charts). We employ a similar idea here, so that our monitoring statistics for the RMCUSUM-V chart become

$$G_t^{*r} \equiv \max_c G_t^{c,r}, \quad t = 1, 2, \dots,$$

for $r \in \mathcal{R}$; and an alarm is raised at t if any G_t^{*r} hits the corresponding control limit, g^{*r} .

Our goal in this section is to develop the RMCUSUM-V chart that can accommodate both the SR-type and DR-type correlations. Note that the monitoring statistic G_t^{*r} is the maximum of p MCUSUM statistics with the same radius, and thus the SR-type correlation can be dealt with simply by searching for the control limit g^{*r} via Phase I of the RMCUSUM-F chart in Figure 17. For this reason, we mainly discuss how to handle the DR-type correlation among the G_t^{*r} s.

If there is no DR-type correlation among the u G_t^{*r} s, finding g^{*r} that yields $u\text{ARL}_0$ for each $r \in \mathcal{R}$ can be good enough to achieve the target ARL_0 of the RMCUSUM-V chart. However, as we saw in similar cases for the SMCUSUM chart (See Pages 40 and 49 in Chapter 3), the chart using the control limits with $u\text{ARL}_0$ may not perform well for the cases where the DR-type correlation actually exists.

For the RMCUSUM-V chart, we adjust u to a corrected value \check{u} by considering the impact of the DR-type correlation, so that the control limits yielding $\check{u}\text{ARL}_0$ for each G_t^{*r} can truly achieve the RMCUSUM-V chart's target ARL_0 . Figure 18 formalizes the procedure for the RMCUSUM-V chart. Note that the RMCUSUM-F chart is a spatial case of the RMCUSUM-V chart with $\check{u} = 1$ and no DR-type correlation.

A remaining question for the RMCUSUM_v chart is, of course, how to find \check{u} values. In order to find \check{u} , we first introduce a measure of the DR-type correlation

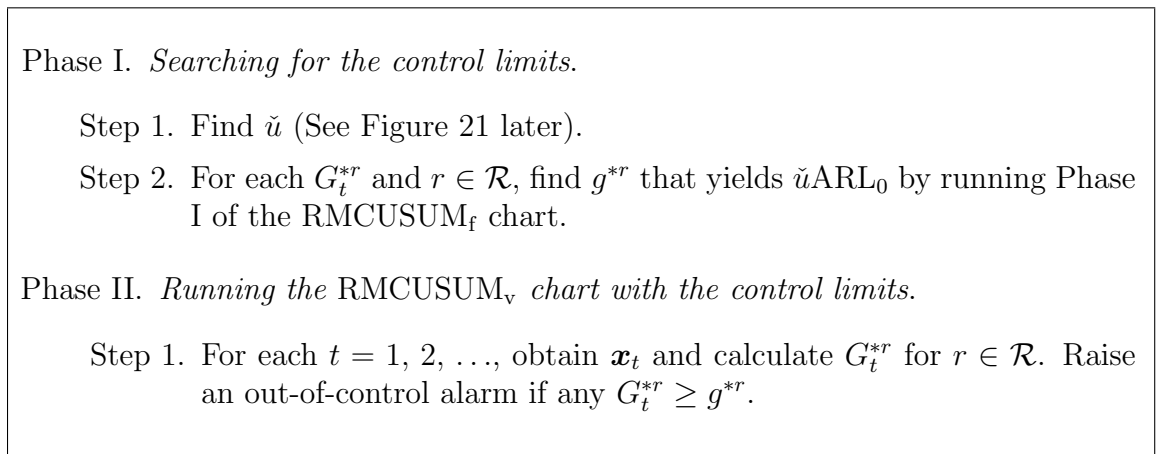


Figure 18: The RMCUSUM_v chart

and a measure of the impact of the DR-type correlation. Then we present a method to estimate \check{u} by using the relationship between the two measures.

Measure of the DR-type correlation: Recall that Ψ_{r_a, r_b} with $r_a = r_b$ is related to the SR-type correlation and Ψ_{r_a, r_b} with $r_a \neq r_b$ is related to the DR-type correlation. Hence, the average value of all components in Ψ_{r_a, r_b} with $r_a \neq r_b$ can be used as a measure of the DR-type correlation. For given p , k , and target ARL_0 , let DRC_ρ denote the average value for a common spatial correlation parameter ρ . Then we use $D_\rho = DRC_\rho / DRC_0$ as a standardized measure of the DR-type correlation for that particular ρ .

Measure of the impact of the DR-type correlation: First, let us assume the $u G_t^{*r}$ s are independent, and suppose we find g^{*r} s that yield $uARL_0$ for the respective G_t^{*r} s by running Phase I of the RMCUSUM-F chart. If we run the RMCUSUM-V chart with the g^{*r} s under the in-control state, we can obtain empirical ARL_0 estimates. Obviously, if the DR-type correlation exists among the G_t^{*r} s, the empirical ARL_0 will likely be different than the target ARL_0 of the RMCUSUM-V chart. Let A_ρ denote the ratio of the empirical ARL_0 to the target ARL_0 for a given ρ (i.e., $A_\rho = \frac{\text{Empirical } ARL_0}{\text{Target } ARL_0}$). Then A_ρ can be used as a measure of the DR-type correlation's impact on the ARL_0 .

We set our correction term $\check{u} = \frac{u}{A_\rho} = u \cdot \frac{\text{Target } ARL_0}{\text{Empirical } ARL_0}$. This means that we set $\check{u} < u$ if the empirical ARL_0 is larger than the target ARL_0 and $\check{u} > u$ if the empirical ARL_0 is smaller than the target. Our method for setting \check{u} is based on an empirical guess, but we will confirm that the \check{u} performs well under various experimental configurations in Section 4.3.

Since obtaining A_ρ for the numerous ρ -values under consideration requires a great deal of simulation time and effort, we present a more-convenient method to estimate A_ρ (and thus \check{u}) by studying the relationship between D_ρ and A_ρ .

Relationship between D_ρ and A_ρ : We plot A_ρ as a function of D_ρ and empirically find a linear relationship between them. More specifically, under each configuration in Table 17, Figure 19 depicts the relationship between D_ρ and A_ρ for Σ_2 with $\rho = 0.0, 0.2, 0.4,$ and 0.7 , where Σ_2 is the same correlation structure defined in Chapter 3. Under the correlation structure Σ_2 , D_ρ is largest (i.e., 1) when $\rho = 0$ and becomes smaller as $\rho > 0$ gets larger. Similarly, A_ρ is largest when $\rho = 0$ (i.e., $D_\rho = 1$) and becomes smaller as $\rho > 0$ gets larger (i.e., D_ρ gets smaller). In Figure

Table 17: Configurations

Configuration	p	k	\mathcal{R}	Target ARL_0
1	7×7	0.5	$\{0, 1, \sqrt{2}, 2\}$	100
2	7×7	0.5	$\{1, \sqrt{2}\}$	100
3	7×7	0.5	$\{\sqrt{2}, 2\}$	100
4	7×7	0.1	$\{\sqrt{2}, 2\}$	100
5	7×7	0.5	$\{\sqrt{2}, 2\}$	370
6	20×20	0.5	$\{\sqrt{2}, 2\}$	100




Figure 19: D_ρ vs A_ρ

19, A_ρ increases linearly as D_ρ increases for the all configurations.

Estimation of A_ρ and \check{u} : Assuming that the linear relationship between D_ρ and A_ρ holds for all configurations, we can estimate A_ρ for various ρ by interpolation. For example, suppose that we first run the RMCUSUM-V chart with ρ_1 and ρ_2 and obtain D_{ρ_1} , D_{ρ_2} , A_{ρ_1} , and A_{ρ_2} . After drawing a line between the two points (D_{ρ_1}, A_{ρ_1}) and (D_{ρ_2}, A_{ρ_2}) as depicted in Figure 20, we can estimate A_ρ for any $\rho_1 \leq \rho \leq \rho_2$ from the following equation,

$$A_\rho \approx \frac{A_{\rho_1} - A_{\rho_2}}{D_{\rho_1} - D_{\rho_2}} (D_\rho - D_{\rho_2}) + A_{\rho_2}, \quad (4.2.3)$$

where D_ρ is the value calculated analytically without the trial-and-error simulation. Since \check{u} can be estimated by u/A_ρ , we can reduce the computational burden when testing various ρ for the RMCUSUM-V chart.

Figure 21 summarizes the procedure of finding \check{u} .




Figure 20: Estimation of A_ρ

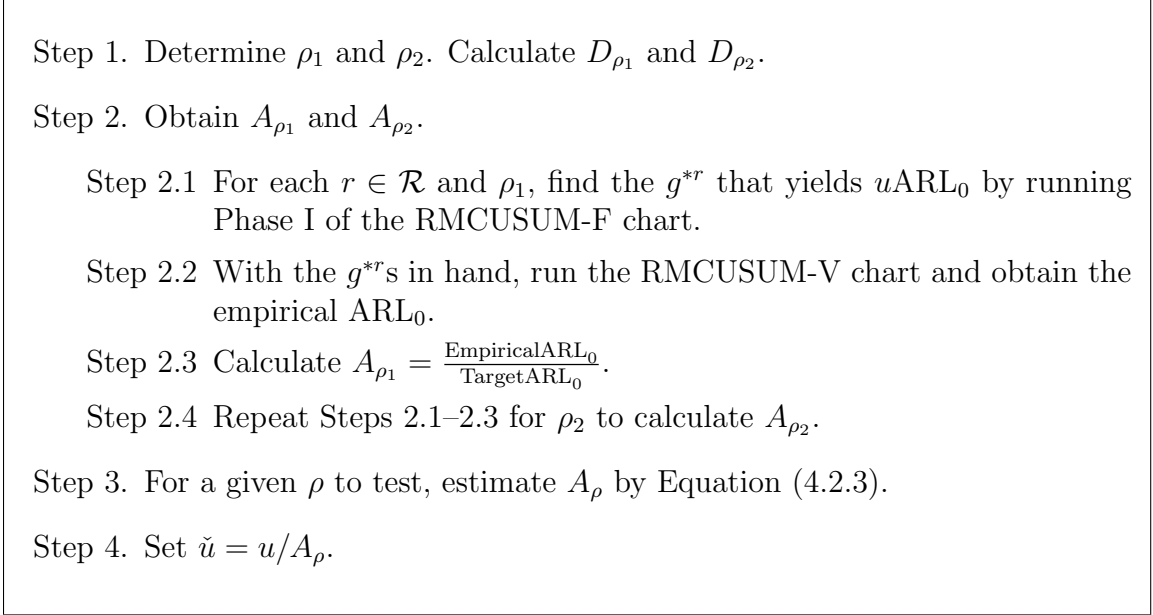


Figure 21: Finding \check{u}

4.3. Experiments

In this section, we test the accuracy of the RMCUSUM charts as well as the robustness of the charts to the underlying distribution of regional data and the correlation structure.

4.3.1 Experimental Setup

For archetypal examples, we consider 7×7 and 20×20 regions for p when the target ARL_0 is set to 100 and 370. For the out-of-control state, we assume a homogeneous shift magnitude with $\delta = 1$.

Outbreak cluster: Outbreak clusters with various r_{out} values are tested when the outbreak center c_{out} is located at the center of the 7×7 or 20×20 grid of the regional system. For example, Figure 22 illustrates outbreak regions in the 7×7 regional system, colored in gray with $c_{\text{out}} = 25$ and $r_{\text{out}} = 0, 1, 2$.

Spatial correlation: We mainly consider two different correlation structures for \mathbf{x}_t . The first correlation structure Σ_1 denotes the $p \times p$ correlation matrix given in Jian et al. [24], where any pair of adjacent regions has a correlation coefficient




Figure 22: Outbreak clusters when $c_{\text{out}} = 25$ (outbreak regions are in gray)

$0 \leq \rho \leq 1$, any two cross-adjacent regions have correlation coefficient $\rho/2$, and any other pairs of regions have no correlation. The second correlation structure Σ_2 denotes the $p \times p$ correlation matrix given in Chapter 3, where any pair of regions has a correlation coefficient that depends on the distance between the regions. Specifically, the (i, j) th element of Σ_2 is $\rho^{\text{dist}(q_i, q_j)}$ for $i, j \in \mathcal{P}$. Under the 7×7 regional system, Figure 23 compares the different correlation settings of Σ_1 and Σ_2 for region 25, which corresponds to the 25th rows of Σ_1 and Σ_2 , respectively. We test the spatial correlation parameter $\rho = 0.0, 0.2, 0.4, 0.7$ for our examples. However, the eigenvalues of Σ_1 approach zero or become negative when $\rho \geq 0.5$, which results in an ill-conditioned Σ_1 . Thus, we restrict $\rho < 0.5$ for Σ_1 .

Marginal distribution: We test two different underlying distributions for the raw data \mathbf{y}_t . One is the multivariate normal distribution with mean vector $\boldsymbol{\mu}$, and the other is the multivariate Poisson distribution with the mean vector whose components are all 5. The reference value k is set to 0.5 for the normal data as is usually used for a classical CUSUM chart, and k is set to 0.1 for the Poisson data as in Kim et al. [25].

Note that a cluster near the edge of the study area can contain a relatively smaller number of regions. In order to keep the number of regions in each cluster the same, we have added dummy regions. For example, we actually generate 9×9 regions to

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	$\rho/2$	ρ	$\rho/2$	0	0
0	0	ρ	1	ρ	0	0
0	0	$\rho/2$	ρ	$\rho/2$	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

$\rho^{\sqrt{18}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{10}}$	ρ^3	$\rho^{\sqrt{10}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{18}}$
$\rho^{\sqrt{13}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{5}}$	ρ^2	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{13}}$
$\rho^{\sqrt{10}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{2}}$	ρ	$\rho^{\sqrt{2}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{10}}$
ρ^3	ρ^2	ρ	1	ρ	ρ^2	ρ^3
$\rho^{\sqrt{10}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{2}}$	ρ	$\rho^{\sqrt{2}}$	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{10}}$
$\rho^{\sqrt{13}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{5}}$	ρ^2	$\rho^{\sqrt{5}}$	$\rho^{\sqrt{8}}$	$\rho^{\sqrt{13}}$
$\rho^{\sqrt{18}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{10}}$	ρ^3	$\rho^{\sqrt{10}}$	$\rho^{\sqrt{13}}$	$\rho^{\sqrt{18}}$

Figure 23: Correlation settings of Σ_1 and Σ_2 for region 25

monitor 7×7 regions with $r = 2$. We conduct 10,000 replications of each experiment to calibrate empirical control limits (i.e., $\tau = 10,000$), and 3,000 replications to obtain empirical ARLs (either ARL_0 or ARL_1). In order to interpolate the lines for the RMCUSUM-V chart, we use $\rho_1 = 0.00$ and $\rho_2 = 0.49$ for Σ_1 , and $\rho_1 = 0.00$ and $\rho_2 = 0.99$ for Σ_2 .

4.3.2 Accuracy of Control Limits

The RMCUSUM charts are designed to be distribution-free. In order to test the accuracy of the control limits of the RMCUSUM charts, we apply the charts with the data following the multivariate normal distribution and the multivariate Poisson distribution under the Σ_2 spatial correlation structure.

Range of the scan radius: Table 18 exhibits the empirical ARL_{0s} of the RMCUSUM charts under various experimental configurations with the multivariate normal data. Note that the ARL_0 values are from the RMCUSUM-F charts when the cardinality of \mathcal{R} is 1, while the ARL_0 values are from the RMCUSUM-V charts when the cardinality of \mathcal{R} is greater than 1. In Table 18, all empirical ARL_{0s} are -3% to 8% off from the target ARL_{0s} for all configurations tested. Unlike the SMCUSUM chart, the RMCUSUM chart does not have any restriction on the range of the scan

Table 18: Empirical ARL_0 from the multivariate normal data

p	\mathcal{R}	Target ARL_0	ρ	Empirical ARL_0
7×7	$\{1\}$	100	0.0	100.65
7×7	$\{1\}$	100	0.4	100.52
7×7	$\{1, \sqrt{2}\}$	100	0.2	104.89
7×7	$\{1, \sqrt{2}\}$	100	0.7	101.46
7×7	$\{\sqrt{2}, 2\}$	100	0.4	97.54
7×7	$\{\sqrt{2}, 2\}$	100	0.7	104.10
7×7	$\{\sqrt{2}, 2\}$	370	0.0	371.74
7×7	$\{\sqrt{2}, 2\}$	370	0.2	393.53
7×7	$\{0, 1, \sqrt{2}, 2\}$	100	0.0	102.68
7×7	$\{0, 1, \sqrt{2}, 2\}$	100	0.4	104.16
20×20	$\{5\}$	100	0.0	100.59
20×20	$\{5\}$	100	0.4	100.44
20×20	$\{\sqrt{2}, 2\}$	100	0.2	99.58
20×20	$\{\sqrt{2}, 2\}$	100	0.7	107.85
20×20	$\{\sqrt{20}, 5\}$	100	0.4	101.36
20×20	$\{\sqrt{20}, 5\}$	370	0.0	371.19

radius. For example, the RMCUSUM chart performs well under the configurations $p = 20 \times 20$ and $\mathcal{R} = \{5\}$ or $\{\sqrt{20}, 5\}$ while the SMCUSUM chart is not able to test those configurations due to the restrictions on its reference value and the range of the scan radius. In addition, the empirical ARL_0 of the SMCUSUM chart decreases as ρ increases because of its assumption on the independence of the MCUSUM statistics. We do not observe such performance deterioration with the RMCUSUM chart. Note that the RMCUSUM-F chart does not require the \tilde{u} estimation procedure and thus can result in less simulation/interpolation error when compared to the RMCUSUM-V chart. This fact is confirmed by the empirical ARL_0 s of the fixed radius cases (within $\pm 1\%$ from the target) in Table 18.

Non-normality: We experiment with multivariate Poisson data in order to test the robustness to non-normality. Table 19 presents the empirical ARL_0 s of the RMCUSUM charts arising from the multivariate Poisson data. The reference value k is set to 0.1, which is employed for analogous Poisson experimentation in Kim et al. [25]. Under all configurations tested, the RMCUSUM chart always gives results

Table 19: Empirical ARL_0 from the multivariate Poisson data

p	\mathcal{R}	Target ARL_0	ρ	Empirical ARL_0
7×7	$\{1\}$	100	0.0	95.04
7×7	$\{1\}$	100	0.4	95.73
7×7	$\{1, \sqrt{2}\}$	100	0.2	96.29
7×7	$\{\sqrt{2}, 2\}$	100	0.0	97.53
7×7	$\{\sqrt{2}, 2\}$	100	0.4	98.96
7×7	$\{\sqrt{2}, 2\}$	370	0.0	354.83
7×7	$\{\sqrt{2}, 2\}$	370	0.2	352.68
7×7	$\{0, 1, \sqrt{2}, 2\}$	100	0.0	94.51
7×7	$\{0, 1, \sqrt{2}, 2\}$	100	0.4	95.67
20×20	$\{5\}$	100	0.2	94.26
20×20	$\{\sqrt{2}, 2\}$	100	0.4	96.12
20×20	$\{\sqrt{2}, 2\}$	370	0.0	354.09

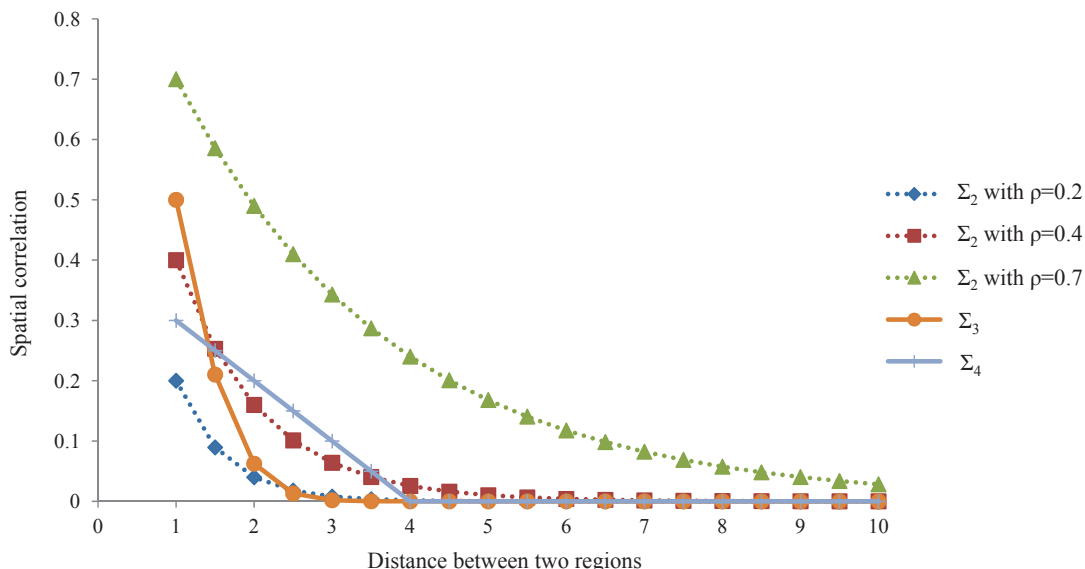


Figure 24: Σ_3 and Σ_4

within -6% of the target ARL_0 s.

Spatial correlation structure: In order to study the robustness to the spatial correlation structure, we test the RMCUSUM charts under the two additional spatial correlation structures, Σ_3 and Σ_4 . The additional correlation structures depend on the distance between the regions. Specifically, the (i, j) th element of Σ_3 is $0.5^{\text{dist}(q_i, q_j)}$ for $i, j \in \mathcal{P}$ while the (i, j) th element of Σ_4 is $0.3 - 0.1(\text{dist}(q_i, q_j) - 1)$. Figure 24

compares Σ_3 and Σ_4 to Σ_2 with $\rho = 0.2, 0.4, 0.7$. When applied to the configuration $p = 7 \times 7$, $\mathcal{R} = \{\sqrt{2}, 2\}$, target $ARL_0 = 100$, and multivariate normal data, the RMCUSUM charts for Σ_3 and Σ_4 provide the empirical ARL_0 s 105.62 and 104.87 respectively.

In all of the configurations we tested, the RMCUSUM chart finds its control limits that deliver ARL_0 s close to the target.

4.3.3 ARL_1 Performance

ARL_1 s are often used as a performance measure for timely detection of outbreaks. In addition to satisfying the pre-specified target ARL_0 , it is also desirable for CUSUM charts to achieve low ARL_1 s.

In Tables 20 and 21, we display the ARL_1 performance of the RMCUSUM charts under Σ_2 when the control limits are searched to yield the target $ARL_0 = 100$. Table 20 reports the ARL_1 s arising from multivariate normal data, while Table 21 reports the ARL_1 s from multivariate Poisson data. In both Tables 20 and 21, the RMCUSUM charts provide lower ARL_1 s as the actual outbreak radius r_{out} gets larger. Further, when the chart's scan radius exactly matches the actual outbreak radius (i.e., $r_{out} = r = 1$), the RMCUSUM chart with the fixed scan radius $\mathcal{R} = \{1\}$ performs better than the chart with variable scan radius $\mathcal{R} = \{0, 1, \sqrt{2}, 2\}$.

All of the ARL_1 s in Table 21 are larger than the corresponding entries in Table 20, which implies that the RMCUSUM charts take more time to detect outbreaks

Table 20: ARL_1 from multivariate normal data

p	ρ	r_{out}	$\mathcal{R} = \{1\}$	$\mathcal{R} = \{0, 1, \sqrt{2}, 2\}$
		0	26.63	14.67
7×7	0.0	1	4.32	4.91
		2	3.39	2.91
		0	21.66	10.57
7×7	0.4	1	4.48	5.38
		2	6.07	3.89

Table 21: ARL_1 from multivariate Poisson data

p	ρ	r_{out}	$\mathcal{R} = \{1\}$	$\mathcal{R} = \{0, 1, \sqrt{2}, 2\}$
7×7	0.0	0	32.7	24.96
		1	8.58	10.44
		2	7.37	6.1
7×7	0.4	0	30.29	20.08
		1	9.13	11.39
		2	12.68	8.52

Table 22: ARL_1 comparison

δ	ρ	r_{out}	RMCUSUM ARL_1 (JMCUSUM ARL_1)			
			$\mathcal{R} = \{0\}$	$\mathcal{R} = \{1\}$	$\mathcal{R} = \{2\}$	$\mathcal{R} = \{0, 1, \sqrt{2}, 2\}$
0.5	0.0	0	41.86 (36.94)	61.59 (62.02)	66.04 (72.16)	44.61 (41.17)
		1	19.00 (21.39)	10.56 (10.41)	13.98 (16.49)	11.59 (10.93)
		2	12.67 (16.58)	7.02 (6.69)	5.25 (4.43)	6.10 (4.93)
	0.2	0	37.37 (33.69)	63.92 (64.77)	71.15 (74.20)	42.33 (37.89)
		1	26.63 (26.39)	12.53 (12.54)	20.30 (23.81)	14.45 (13.78)
		2	20.23 (22.77)	10.95 (10.86)	6.99 (6.42)	8.22 (7.18)
0.4	0	20.88 (21.28)	43.58 (47.86)	55.42 (66.20)	23.72 (24.22)	
	1	16.39 (17.54)	9.30 (9.02)	15.57 (18.66)	10.96 (10.39)	
	2	15.09 (16.66)	10.20 (10.00)	6.32 (5.63)	7.52 (6.41)	
1.0	0.0	0	13.17 (13.08)	26.63 (47.75)	36.40 (66.68)	14.67 (14.60)
		1	7.46 (7.46)	4.32 (3.36)	5.72 (7.57)	4.91 (3.66)
		2	5.93 (5.95)	3.39 (2.24)	2.50 (1.45)	2.91 (1.67)
	0.2	0	11.54 (11.66)	27.16 (44.04)	37.92 (67.81)	13.34 (13.40)
		1	9.54 (9.56)	4.94 (4.08)	7.89 (11.15)	5.76 (4.60)
		2	8.40 (8.24)	4.93 (3.74)	3.16 (2.11)	3.69 (2.42)
0.4	0	7.17 (6.91)	15.06 (28.31)	22.05 (54.09)	8.35 (8.03)	
	1	6.53 (6.13)	3.91 (2.92)	6.32 (7.92)	4.58 (3.39)	
	2	6.51 (5.95)	4.63 (3.40)	2.89 (1.84)	3.39 (2.20)	

in the case of multivariate Poisson data than with multivariate normal data. This is because the multivariate Poisson distribution we test is asymmetric and generates negative values for the standardized \mathbf{x}_t more frequently than does the multivariate normal distribution; and the reference value k we use for the multivariate Poisson data is different than the k for the multivariate normal data.

Table 22 compares the ARL_1 s of the RMCUSUM chart to the ARL_1 s of the JMCUSUM chart under Σ_2 . For all configurations in Table 22, the target ARL_0

is set to 100. The RMCUSUM charts provide significantly lower ARL_1 s than the JMCUSUM charts for the fixed radius cases $r > r_{out}$. With the variable radius, the JMCUSUM charts produce slightly lower ARL_1 s than the RMCUSUM charts. Overall, ARL_1 performances of the RMCUSUM and JMCUSUM charts are similar.

4.4. Conclusions

Motivated by the JMCUSUM and SMCUSUM charts, we wish to develop a new method that can search for the control limits of correlated MCUSUM charts. The new method is named the RMCUSUM chart, and is designed to be distribution-free for the underlying data, restriction-free for the range of scan radius, and robust to spatial correlation structures.

The RMCUSUM chart determines its control limits via simple simulation and interpolation. Our search method requires the variance-covariance matrix of the data and the generation of independent standard normal observations, but does not require modeling of the underlying process.

Control limits for the RMCUSUM charts are accurately determined under various problem configurations. We test the RMCUSUM charts with four different spatial correlation structures and two underlying distributions of data that have been considered in the biosurveillance literature. We find a linear relationship between the suggested measure of DR-type correlation and its impact on empirical ARL_0 s, and we develop a technique to reduce the computational burden for testing the RMCUSUM charts.

While maintaining satisfactory accuracy of the control limits, the RMCUSUM charts produce comparable ARL_1 s to those from the JMCUSUM charts; but the RMCUSUM charts perform better when we use a fixed scan radius larger than the actual outbreak radius.

Chapter V

CONTRIBUTIONS

The epidemic simulation model discussed in this thesis reports valid estimates of illness attack rates and economic costs associated with different combinations of mitigation strategies. Our research shows that delays in vaccination can greatly reduce the effectiveness of the vaccination strategy, and that reactive combination strategies including practical school closure can substantially reduce total costs associated with influenza morbidity and mortality. We also provide the information needed for users to build a simulation model for disease propagation, to use the model to better understand diseases, and to identify appropriate intervention strategies.

In order to decide when and where to apply the mitigation strategies, the detection of an outbreak should be accomplished at an early time. In terms of statistical approaches for outbreak detection in the presence of spatial correlation, we develop spatiotemporal monitoring statistics for two MCUSUM charts, named SMCUSUM and RMCUSUM.

From the SMCUSUM chart, we study how the outbreak coverage and spatial correlation impact the charts's performance, and find that higher spatial correlation does not always yield faster detection but often facilitates accurate identification of outbreaks. When assuming the independence among the MCUSUM statistics and confining the scan radius to a certain range, we approximate the control limits analytically so that we can completely avoid the computational burden for trial-and-error simulations to find the control limit. The approximation method is effective especially when the number of monitoring regions is large.

When severe correlations exist among the MCUSUM statistics or when the scan

radius should not be restricted, one can use the RMCUSUM chart. A critical strength of the control limit approximation method for the RMCUSUM chart is that we can use such a chart even when the data distribution is unknown — without modeling of the complicated underlying process. Based on our experimental results, we believe that the RMCUSUM chart is robust enough to handle a variety of data distributions and spatial correlation structures.

By using the epidemic simulation model and the MCUSUM charts suggested in this thesis, we expect to help public health authorities, practitioners, researchers, and decision makers when preparing or reacting to future outbreaks of pandemics.

APPENDIX A

Table 23: Average age-stratified overall illness attack rates (%) and cost estimates

Intervention ¹	Average Overall Illness Attack Rate (%)						Cost (US\$m)
	Pre-schoolers	School children	Young adults	Adults	Older Adults	Overall	
No Intervention	29.5	55.9	40.8	14.3	11	34.1	81.1
Vaccination Only, Low Efficacy, 35% Coverage							
Pre-vaccination	21.2	47.9	30.2	9.5	7.3	26.1	71.1
Reactive Vaccination ² :							
Initial Dly: None Addl Dly: Yes	24	51	33.8	11	8.5	28.8	77.7
Initial Dly: 30 Day Addl Dly: Yes	24.6	51.6	34.6	11.4	8.8	29.5	79.3
Initial Dly: 60 Day Addl Dly: Yes	27.4	54.2	38.3	13.1	10.1	32.2	86
Initial Dly: None Addl Dly: No	24	51	33.8	11	8.4	28.8	77.7
Initial Dly: 30 Day Addl Dly: No	24.2	51.3	34	11.1	8.5	29	78.1
Initial Dly: 60 Day Addl Dly: No	25.6	52.6	36.3	12.2	9.3	30.7	82.2
Vaccination Only, Low Efficacy, 70% Coverage							
Pre-vaccination	8.8	25.3	13.1	3.6	2.8	12	47
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	17.8	43.2	25.4	7.7	5.9	22.4	71.6
Initial Dly: 30 Day Addl Dly: Yes	19.2	45.1	27.7	8.6	6.6	24.1	75.7
Initial Dly: 60 Day Addl Dly: Yes	25.4	52.2	36	12.1	9.3	30.4	89.4
Initial Dly: None Addl Dly: No	17.4	42.9	25.1	7.6	5.8	22.2	71.1
Initial Dly: 30 Day Addl Dly: No	18.1	43.5	25.8	7.9	6	22.7	72.4
Initial Dly: 60 Day Addl Dly: No	22	48.2	31.8	10.3	7.8	27.1	83
Vaccination Only, Moderate Efficacy, 35% Coverage							
Pre-vaccination	14.4	37.5	21.1	6.3	4.7	18.8	53.7
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	18	43.4	26	8	6.1	22.8	63.1
Initial Dly: 30 Day Addl Dly: Yes	19.6	45.5	28.4	8.9	6.8	24.6	67.5
Initial Dly: 60 Day Addl Dly: Yes	25.7	52.5	36.5	12.3	9.4	30.8	82.5
Initial Dly: None Addl Dly: No	18	43.2	25.8	7.9	6	22.6	62.8
Initial Dly: 30 Day Addl Dly: No	18.3	43.7	26.4	8.1	6.2	23	63.7
Initial Dly: 60 Day Addl Dly: No	22.2	48.3	32.1	10.3	8	27.3	74.1
Vaccination Only, Moderate Efficacy, 70% Coverage							
Pre-vaccination	0.1	0.4	0.2	0.04	0.03	0.2	19.3
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	3.3	10.1	5	1.3	1	4.6	29.7
Initial Dly: 30 Day Addl Dly: Yes	9.7	26.8	15	4.3	3.3	13.3	50.2
Initial Dly: 60 Day Addl Dly: Yes	22.3	48.3	32.6	10.6	8.1	27.6	83
Initial Dly: None Addl Dly: No	1.5	4.7	2.3	0.6	0.5	2.2	25.6
Initial Dly: 30 Day Addl Dly: No	5.7	17.1	8.9	2.4	1.9	8.1	39.5
Initial Dly: 60 Day Addl Dly: No	17.5	40.5	26.7	8.3	6.3	22.6	74
Antivirals (10% Cov.) Only	26	53	37.3	12.7	9.6	31.3	75.9
SC/SD, Only	22.5	43.6	27.7	8.6	6.7	24	125
Antivirals (10% Cov.)+SC/SD	7.1	18.8	10.2	2.8	2.2	9.2	48

¹ We use the abbreviation SC/SD for School closure/Social distancing intervention.

² In reactive vaccination scenarios, two types of supply-chain delays that can affect vaccination programs are considered: an initial delay (initial dly) in program implementation of 0, 30, or 60 days; and additional delays (addl dly) after initiation of the program, such that vaccine doses become available in three equal batches, two weeks apart (rather than in one batch).

Table 23 (Continued)

Intervention	Average Overall Illness Attack Rate (%)					Overall	Cost (US\$m)
	Pre-schoolers	School children	Young adults	Adults	Older Adults		
Vaccination (Low Efficacy, 35% Coverage) + Antivirals (10% Coverage)							
Pre-vaccination	14	37.8	22	6.5	5	19.3	56.4
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	19.6	46.9	29.3	9.2	7	25.3	70.8
Initial Dly: 30 Day Addl Dly: Yes	20.1	47.4	29.9	9.4	7.1	25.7	71.8
Initial Dly: 60 Day Addl Dly: Yes	21.4	48.8	31.6	10.1	7.7	27.1	75
Initial Dly: None Addl Dly: No	19.5	46.7	29.2	9.1	6.9	25.2	70.6
Initial Dly: 30 Day Addl Dly: No	19.7	47	29.5	9.2	7	25.4	71.1
Initial Dly: 60 Day Addl Dly: No	20.6	47.9	30.5	9.6	7.3	26.2	72.9
Vaccination (Low Efficacy, 70% Coverage) + Antivirals (10% Coverage)							
Pre-vaccination	2	7.3	3.6	0.9	0.7	3.3	28.3
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	12.6	35.9	19.8	5.7	4.4	17.7	62
Initial Dly: 30 Day Addl Dly: Yes	13.1	37.2	20.7	6	4.6	18.4	63.9
Initial Dly: 60 Day Addl Dly: Yes	16.5	42.3	25.2	7.6	5.8	22	72.4
Initial Dly: None Addl Dly: No	12.2	34.9	19.4	5.6	4.2	17.3	61.1
Initial Dly: 30 Day Addl Dly: No	12.6	36.2	20	5.8	4.4	17.9	62.5
Initial Dly: 60 Day Addl Dly: No	14.6	39.6	22.5	6.6	5.1	19.9	67.4
Vaccination (Moderate Efficacy, 35% Coverage) + Antivirals (10% Coverage)							
Pre-vaccination	1.2	4.8	2.3	0.6	0.5	2.1	16.1
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	6.2	21.1	11.1	3	2.3	10	34.3
Initial Dly: 30 Day Addl Dly: Yes	10.9	32.2	17.8	5.1	3.9	15.8	48.2
Initial Dly: 60 Day Addl Dly: Yes	15.4	40.4	23.8	7.1	5.4	20.8	60.1
Initial Dly: None Addl Dly: No	4.8	17.4	9	2.4	1.8	8.1	30.1
Initial Dly: 30 Day Addl Dly: No	8.2	25.9	13.9	3.9	2.9	12.4	40.2
Initial Dly: 60 Day Addl Dly: No	13.3	36.9	21	6.1	4.7	18.6	54.7
Vaccination (Moderate Efficacy, 70% Coverage) + Antivirals (10% Coverage)							
Pre-vaccination	0.03	0.14	0.05	0.01	0.01	0.05	20.7
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	0.7	2.7	1.2	0.3	0.2	1.2	23.3
Initial Dly: 30 Day Addl Dly: Yes	2.6	1.6	4.8	1.3	1	4.4	30.9
Initial Dly: 60 Day Addl Dly: Yes	8	25	13.7	3.8	2.9	12.2	49.1
Initial Dly: None Addl Dly: No	0.4	1.5	0.6	0.2	0.1	0.6	22
Initial Dly: 30 Day Addl Dly: No	1.4	5.4	2.6	0.7	0.5	2.4	26.2
Initial Dly: 60 Day Addl Dly: No	3.9	14	7.4	2	1.5	6.6	36.1
Vaccination (Low Efficacy, 35% Coverage) + SC/SD							
Pre-vaccination	10.9	26.3	13.9	3.9	3.1	12.7	69.9
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	15.5	34.7	19.6	5.7	4.5	17.5	95.7
Initial Dly: 30 Day Addl Dly: Yes	16.3	35.9	20.4	6	4.7	18.3	99
Initial Dly: 60 Day Addl Dly: Yes	17.5	37.7	22.1	6.6	5.1	19.6	108.8
Initial Dly: None Addl Dly: No	15.3	34.3	19.3	5.6	4.4	17.3	93.6
Initial Dly: 30 Day Addl Dly: No	15.8	35.2	19.9	5.8	4.5	17.8	96.5
Initial Dly: 60 Day Addl Dly: No	16.7	36.3	20.9	6.2	4.8	18.6	101.9
Vaccination (Low Efficacy, 70% Coverage) + SC/SD							
Pre-vaccination	0.5	1.6	0.7	0.2	0.1	0.7	22
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	6.1	16	8.1	2.2	1.7	7.5	53.1
Initial Dly: 30 Day Addl Dly: Yes	9.2	23.3	12	3.3	2.6	11	70.6
Initial Dly: 60 Day Addl Dly: Yes	13.3	30.7	17.1	4.9	3.9	15.4	96.6
Initial Dly: None Addl Dly: No	4.8	12.7	6.3	1.7	1.3	5.9	46
Initial Dly: 30 Day Addl Dly: No	7.9	10.3	10.3	2.8	2.2	9.5	63
Initial Dly: 60 Day Addl Dly: No	11.3	27.3	14.6	4.1	3.2	13.3	82.6

Table 23 (Continued)

Intervention	Average Overall Illness Attack Rate (%)						Cost (US\$m)
	Pre-schoolers	School children	Young adults	Adults	Older Adults	Overall	
Vaccination (Moderate Efficacy, 35% Coverage) + SC/SD							
Pre-vaccination	1.8	5	2.4	0.6	0.5	2.3	19.6
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	7.1	18	9.3	2.5	2	8.5	49.4
Initial Dly: 30 Day Addl Dly: Yes	13.5	30.7	17.1	4.9	3.9	15.4	87.3
Initial Dly: 60 Day Addl Dly: Yes	15.8	35.5	19.9	5.8	4.5	17.9	95.7
Initial Dly: None Addl Dly: No	5.6	14.4	7.4	2	1.6	6.8	41.6
Initial Dly: 30 Day Addl Dly: No	8.4	10.8	10.8	3	2.3	9.9	56.3
Initial Dly: 60 Day Addl Dly: No	11.6	27.3	14.8	4.2	3.3	13.4	74.7
Vaccination (Moderate Efficacy, 70% Coverage) + SC/SD							
Pre-vaccination	0.03	0.1	0.04	0.01	0.01	0.04	19.1
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	0.5	1.6	0.7	0.2	0.1	0.7	22
Initial Dly: 30 Day Addl Dly: Yes	2.6	7	3.5	0.9	0.7	3.2	34.7
Initial Dly: 60 Day Addl Dly: Yes	8.4	19.9	10.9	3	2.4	9.8	69.1
Initial Dly: None Addl Dly: No	0.15	0.48	0.19	0.05	0.04	0.19	19.7
Initial Dly: 30 Day Addl Dly: No	1.21	3.29	1.58	0.42	0.32	1.49	25.9
Initial Dly: 60 Day Addl Dly: No	5.19	13.11	7.01	1.88	1.48	6.35	51.2
Vaccination (Low Efficacy, 35% Coverage) + Antivirals (10% Coverage) + SC/SD							
Pre-vaccination	0.7	2.4	1.1	0.3	0.2	1	15.9
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	3.1	9.8	4.8	1.3	1	4.5	32.2
Initial Dly: 30 Day Addl Dly: Yes	3.3	10.7	5.4	1.4	1.1	4.9	34.2
Initial Dly: 60 Day Addl Dly: Yes	3.6	11.7	5.6	1.5	1.2	5.4	36.8
Initial Dly: None Addl Dly: No	2.6	8.4	4.2	1.1	0.8	3.9	29.2
Initial Dly: 30 Day Addl Dly: No	3.1	10	5	1.3	1	4.6	32.6
Initial Dly: 60 Day Addl Dly: No	3.2	10.4	5.2	1.4	1.1	4.8	33.8
Vaccination (Low Efficacy, 70% Coverage) + Antivirals (10% Coverage) + SC/SD							
Pre-vaccination	0.2	0.5	0.2	0.1	0.04	0.2	21.3
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	1.8	5.9	2.7	0.7	0.6	2.6	32
Initial Dly: 30 Day Addl Dly: Yes	2.3	7.1	3.4	0.9	0.7	3.2	35.2
Initial Dly: 60 Day Addl Dly: Yes	3.2	10	4.9	1.3	1	4.6	41.7
Initial Dly: None Addl Dly: No	1.2	3.9	1.9	0.5	0.4	1.8	28.4
Initial Dly: 30 Day Addl Dly: No	2	6.4	3.1	0.8	0.6	2.9	33.6
Initial Dly: 60 Day Addl Dly: No	2.7	8.4	4.1	1.1	0.8	3.8	37.8
Vaccination (Moderate Efficacy, 35% Coverage) + Antivirals (10% Coverage) + SC/SD							
Pre-vaccination	0.1	0.6	0.2	0.1	0.04	0.2	11.9
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	0.5	1.9	0.9	0.2	0.2	0.8	14.9
Initial Dly: 30 Day Addl Dly: Yes	1.1	3.7	1.7	0.5	0.3	1.6	18.6
Initial Dly: 60 Day Addl Dly: Yes	1.5	5.3	2.5	0.6	0.5	2.4	22
Initial Dly: None Addl Dly: No	0.3	1.1	0.5	0.1	0.1	0.5	13.1
Initial Dly: 30 Day Addl Dly: No	0.8	2.8	1.3	0.3	0.3	1.2	16.6
Initial Dly: 60 Day Addl Dly: No	1.2	4.4	2.1	0.5	0.4	2	20.2
Vaccination (Moderate Efficacy, 70% Coverage) + Antivirals (10% Coverage) + SC/SD							
Pre-vaccination	0.02	0.05	0.02	0.01	0	0.02	20.6
Reactive Vaccination:							
Initial Dly: None Addl Dly: Yes	0.2	0.6	0.2	0.1	0.1	0.2	21.6
Initial Dly: 30 Day Addl Dly: Yes	0.5	1.6	0.7	0.2	0.1	0.7	23.8
Initial Dly: 60 Day Addl Dly: Yes	0.9	3.2	1.5	0.4	0.3	1.4	27.4
Initial Dly: None Addl Dly: No	0.1	0.3	0.1	0.03	0.02	0.1	20.1
Initial Dly: 30 Day Addl Dly: No	0.3	1.1	0.5	0.1	0.1	0.5	22.8
Initial Dly: 60 Day Addl Dly: No	0.8	2.7	1.3	0.3	0.3	1.2	26.1

APPENDIX B

Additional experimental results are provided for $p = 7 \times 7$, 20×20 , and $ARL_0 = 100, 370$. Based on the results from Chapter 3, we expect similar performances with Σ_1 and Σ_2 , and thus run the additional experiments only with Σ_2 . For the out-of-control state when $p = 20 \times 20$, we have the outbreak center $c_{\text{out}} = 210$ (i.e., at coordinate $q_c = (11, 10)$) with $r_{\text{out}} = 0, 1, 2$, and assume a homogeneous shift magnitude $\delta = 1$.

For the JMCUSUM-F chart with $p = 20 \times 20$, we provide a series of figures corresponding to Figures 12, 13, and 14, respectively. We only test the target values $ARL_0 = 100$ and 370 here; we expect that the tendency of the ratio (the empirical ARL_0 to the target ARL_0) with a larger target $ARL_0 = 700$ would not be significantly different from the tendency observed in Section 3.3.2. Figure 25 depicts the impact of r on the ARL approximation accuracy for a single MCUSUM chart. Figure 26 shows how the ratio changes for a classical CUSUM chart as k increases. Figure 27 gives the performance of the JMCUSUM-F chart. In order to keep $k \leq 2$ for all ρ values tested, we consider $r \in \{0, 1, \sqrt{2}, 2\}$.

The detailed experimental results are presented in Tables 24, 25, 26, and 27. For the fixed scan radius case, empirical ARL_0 s of the JMCUSUM-F chart with analytical control limits are reported for some selected scan radius options in Table 24. Table 25 compares analytical ARL_1^{JF} s to empirical ARL_1^{JF} s with analytical control limits. For the variable scan radius case, Table 26 presents empirical ARL_0 s of the SMCUSUM chart with analytical control limits, and Table 27 compares empirical ARL_1 s of the JMCUSUM-V chart and SMCUSUM chart under various experimental configurations.

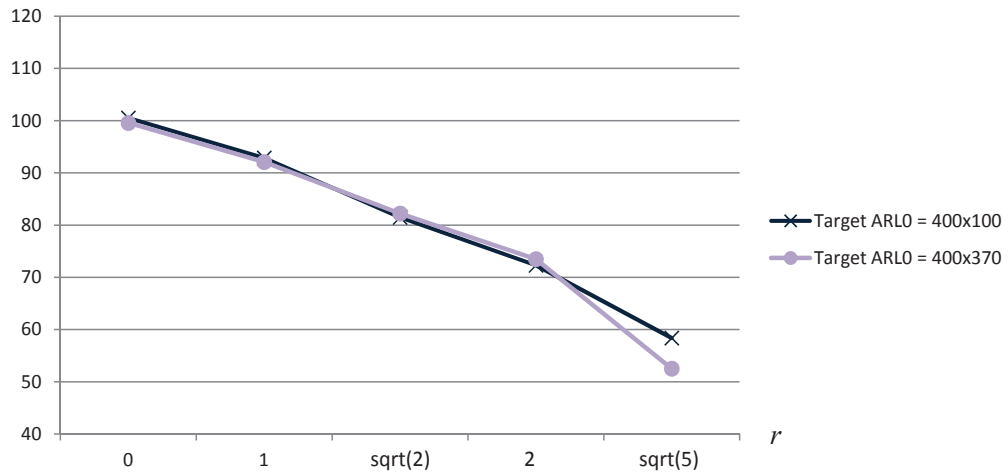


Figure 25: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus r for a single MCUSUM

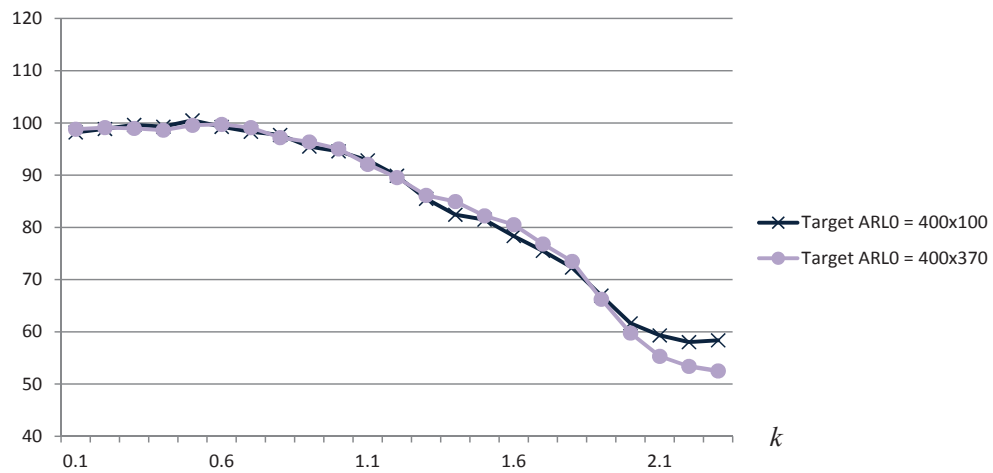


Figure 26: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0}$ (%) versus k for a single MCUSUM




Figure 27: $\frac{\text{Empirical ARL}_0}{\text{Target ARL}_0} (\%)$ versus r for the JMCUSUM-F chart

Table 24: Empirical ARL_0 of the JMCUSUM-F chart with analytical h^{**}

p	Target ARL_0	ρ	$r = 0$	$r = 1$	$r = 2$
7×7	370	0.0	377.42	355.96	298.78
		0.2	385.42	358.85	326.63
		0.4	376.32	352.58	315.08
		0.7	353.37	292.37	239.32
20×20	100	0.0	113.60	97.96	81.55
		0.2	112.76	99.82	90.74
		0.4	107.32	97.57	88.49
		0.7	100.02	82.38	64.75
20×20	370	0.0	383.59	358.96	288.24
		0.2	386.37	356.28	325.12
		0.4	371.55	349.61	323.49
		0.7	360.43	302.55	233.97

Table 25: Empirical (analytical) ARL_1^{JF} with analytical h^{**}

p	Target ARL_0	ρ	$r = r_{out} = 0$	$r = r_{out} = 1$	$r = r_{out} = 2$
7×7	370	0.0	16.18 (16.23)	3.92 (3.89)	1.70 (1.64)
		0.2	14.32 (14.32)	4.76 (4.65)	2.50 (2.43)
		0.4	10.68 (10.65)	4.13 (4.07)	2.57 (2.47)
		0.7	5.16 (5.08)	2.22 (2.10)	1.50 (1.39)
20×20	100	0.0	17.57 (17.81)	4.23 (4.21)	1.82 (1.76)
		0.2	15.43 (15.69)	5.04 (5.04)	2.70 (2.62)
		0.4	11.51 (11.63)	4.45 (4.40)	2.74 (2.66)
		0.7	5.56 (5.50)	2.37 (2.26)	1.60 (1.49)
20×20	370	0.0	20.27 (20.42)	4.79 (4.73)	2.07 (1.97)
		0.2	18.04 (17.96)	5.77 (5.68)	3.04 (2.93)
		0.4	13.30 (13.25)	5.06 (4.95)	3.11 (2.97)
		0.7	6.28 (6.21)	2.63 (2.52)	1.80 (1.65)

Table 26: Empirical ARL_0 for the SMCUSUM chart with analytical h^{*r}

p	Target ARL_0	ρ	ARL_0
7×7	370	0.0	383.18
		0.2	368.12
		0.4	349.58
		0.7	326.57
20×20	100	0.0	107.45
		0.2	100.38
		0.4	95.52
		0.7	90.30
20×20	370	0.0	367.94
		0.2	348.87
		0.4	337.64
		0.7	325.26

Table 27: Empirical ARL_1^{JV} (ARL_1^S)

p	Target ARL_0	ρ	$r_{out} = 0$	$r_{out} = 1$	$r_{out} = 2$
7×7	370	0.0	17.75 (17.94)	4.25 (4.29)	1.90 (1.92)
		0.2	16.64 (16.67)	5.28 (5.39)	2.89 (2.85)
		0.4	12.54 (12.69)	4.80 (4.87)	3.02 (2.98)
		0.7	5.64 (5.70)	2.53 (2.45)	1.49 (1.68)
20×20	100	0.0	18.37 (18.57)	4.91 (4.63)	2.03 (2.05)
		0.2	16.96 (16.98)	5.65 (5.75)	3.07 (3.05)
		0.4	13.05 (13.15)	5.11 (5.11)	3.20 (3.15)
		0.7	6.01 (6.03)	2.65 (2.59)	1.59 (1.79)
20×20	370	0.0	22.07 (22.07)	5.13 (5.16)	2.27 (2.28)
		0.2	20.35 (20.36)	6.42 (6.47)	3.41 (3.37)
		0.4	15.24 (15.31)	5.71 (5.74)	3.54 (3.50)
		0.7	6.79 (6.82)	2.93 (2.87)	1.77 (1.98)

APPENDIX C

In this Appendix, we provide the detailed derivation of Equation (4.2.2).

First, note that

$$\mathbf{S}(t) = \begin{bmatrix} S_1(t) \\ S_2(t) \\ \vdots \\ S_p(t) \end{bmatrix} = -\mathbf{K}t + \sum_{\ell=1}^t \boldsymbol{\alpha}(\ell)$$

where \sum denotes the componentwise summation of vectors.

If Γ is positive-definite, there exist the matrices Λ and Q that satisfy $\Gamma = Q\Lambda Q' = Q\Lambda^{\frac{1}{2}}(Q\Lambda^{\frac{1}{2}})'$ [47]. From the eigen decomposition, Λ is the $p \times p$ diagonal matrix whose entries are the eigenvalues of Γ and Q is the $p \times p$ orthogonal matrix (i.e., $Q' = Q^{-1}$) whose columns are the eigenvectors of Γ .

From the similar idea employed in Kim and Dieker [26], for $\mathbf{V}(t) \equiv \Lambda^{-\frac{1}{2}}Q'\boldsymbol{\alpha}(\ell)$,

$$\mathbb{E}[\mathbf{V}(\ell)] = \mathbf{0}$$

in the in-control state and

$$\text{Var}[\mathbf{V}(\ell)] = \text{Var}\left[\Lambda^{-\frac{1}{2}}Q'\boldsymbol{\alpha}(\ell)\right] = \Lambda^{-\frac{1}{2}}Q'\Gamma Q\Lambda^{-\frac{1}{2}} = \mathbf{I}_p.$$

Then, by Equation (4.1.2),

$$\sum_{\ell=1}^t \mathbf{V}(\ell) \stackrel{\text{D}}{\approx} \mathbf{W}(t).$$

Therefore, where we set $\Gamma^{\frac{1}{2}} = Q\Lambda^{\frac{1}{2}}$,

$$\begin{aligned} \mathbf{S}(t) &= -\mathbf{K}t + \sum_{\ell=1}^t \boldsymbol{\alpha}(\ell) \\ &= -\mathbf{K}t + Q\Lambda^{\frac{1}{2}} \sum_{\ell=1}^t \mathbf{V}(\ell) \\ &\stackrel{\text{D}}{\approx} -\mathbf{K}t + \Gamma^{\frac{1}{2}}\mathbf{W}(t). \end{aligned}$$

REFERENCES

- [1] ALEMAN, D., WIBISONO, T., and SCHWARTZ, B., “Accounting for individual behaviors in a pandemic disease spread model,” in *Proceedings of the 2009 Winter Simulation Conference* (ROSSETTI, M. D., HILL, R. R., JOHANSSON, B., DUNKIN, A., and INGALLS, R. G., eds.), (Austin, TX , USA), IEEE, December 2009.
- [2] ANDRADÓTTIR, S., CHIU, W., GOLDSMAN, D., LEE, M. L., TSUI, K. L., SANDER, B., FISMAN, D. N., and NIZAM, A., “Reactive strategies for containing developing outbreaks of pandemic influenza,” *BMC Public Health*, vol. 11, no. Suppl 1, p. S1, 2011.
- [3] BASTA, N. E., CHAO, D. L., HALLORAN, M. E., MATRAJT, L., and LONGINI JR, I. M., “Strategies for pandemic and seasonal influenza vaccination of schoolchildren in the united states,” *American Journal of Epidemiology*, vol. 170, no. 6, pp. 679–686, 2009.
- [4] BILLINGSLEY, P., *Convergence of probability measures*. Wiley, New York, 1968.
- [5] BOWLBY, G., “Provincial drop-out rates-trends and consequences,” *Education Matters*, vol. 2, no. 4, pp. 81–004–XIE, <http://www.statcan.gc.ca/pub/81-004-x/2005004/8984-eng.htm#b>, 2005.
- [6] BUREAU OF LABOR STATISTICS, “Consumer price index,” http://www.bls.gov/cpi/cpi_dr.htm (Accessed November 29, 2009).
- [7] CENTERS FOR DISEASE CONTROL AND PREVENTION, “2009 H1N1 flu situation update,” <http://www.cdc.gov/H1n1flu/update.htm> (Accessed October 26, 2009).

- [8] CENTERS FOR DISEASE CONTROL AND PREVENTION, “2009 H1N1 vaccination recommendations,” <http://www.cdc.gov/h1n1flu/vaccination/acip.htm> (Accessed December 9, 2009).
- [9] CENTERS FOR DISEASE CONTROL AND PREVENTION, “2009 H1N1 vaccine doses allocated, ordered, and shipped by project area,” <http://www.cdc.gov/h1n1flu/vaccination/updates/102309.htm> (Accessed October 26, 2009).
- [10] CROSIER, R. B., “Multivariate generalizations of cumulative sum quality-control schemes,” *Technometrics*, vol. 30, pp. 291–303, 1988.
- [11] DAVID, H. A., “On the application to statistics of an elementary theorem in probability,” *Biometrika*, vol. 43, no. 1, pp. 85–91, 1956.
- [12] FERGUSON, N. M., CUMMINGS, D. A. T., CAUCHEMEZ, S., FRASER, C., RILEY, S., MEEYAI, A., IAMSIRITHAWORN, S., and BURKE, D. S., “Strategies for containing an emerging influenza pandemic in southeast asia,” *Nature*, vol. 437, no. 7056, pp. 209–214, 2005.
- [13] FRASER, C., DONNELLY, C. A., CAUCHEMEZ, S., HANAGE, W. P., VAN KERKHOVE, M. D., HOLLINGSWORTH, T. D., GRIFFIN, J., BAGGALLEY, R. F., JENKINS, H. E., LYONS, E. J., JOMBART, T., HINSLEY, W. R., GRASSLY, N. C., BALLOUX, F., GHANI, A. C., FERGUSON, N. M., RAMBAUT, A., PYBUS, O. G., LOPEZ-GATELL, H., ALPUCHE-ARANDA, C. M., CHAPELA, I. B., ZAVALA, E. P., GUEVARA, D. M. E., CHECCHI, F., GARCIA, E., and HUGONNET, S., “Pandemic potential of a strain of influenza A (H1N1): early findings,” *Science*, vol. 324, no. 5934, pp. 1557–1561, 2009.
- [14] FRICKER, R. D., HEGLER, B. L., and DUNFEE, D. A., “Comparing syndromic surveillance detection methods: EARS’ versus a CUSUM-based methodology,” *Statistics in Medicine*, vol. 27, no. 17, pp. 3407–3429, 2008.

- [15] GANGNON, R. E. and CLAYTON, M. K., “Bayesian detection and modeling of spatial disease clustering,” *Biometrics*, vol. 56, pp. 922–935, 2000.
- [16] GANGNON, R. E. and CLAYTON, M. K., “A weighted average likelihood ratio test for spatial clustering of disease,” *Statistics in Medicine*, vol. 20, pp. 2977–2987, 2001.
- [17] GOJOVIC, M. Z., SANDER, B., FISMAN, D., KRAHN, M. D., and BAUCH, C. T., “Modelling mitigation strategies for pandemic (H1N1) 2009,” *Canadian Medical Association Journal*, vol. 181, no. 10, pp. 673–680, 2009.
- [18] HALLORAN, M. E., FERGUSON, N. M., EUBANK, S., LONGINI, I. M., CUMMINGS, D. A. T., LEWIS, B., XU, S., FRASER, C., VULLIKANTI, A., GERMAN, T. C., WAGENER, D., BECKMAN, R., KADAU, K., BARRETT, C., MACKEN, C. A., BURKE, D. S., and COOLEY, P., “Modeling targeted layered containment of an influenza pandemic in the united states,” *Proceedings of the National Academy of Sciences*, vol. 105, no. 12, pp. 4639–4644, 2008.
- [19] HAN, S. W., TSUI, K.-L., ARIYAJUNYA, B., and KIM, S. B., “A comparison of CUSUM, EWMA, and temporal scan statistics for detection of increases in Poisson rates,” *Quality and Reliability Engineering International*, vol. 26, pp. 279–289, 2010.
- [20] HANCOCK, K., VEGUILLA, V., LU, X., ZHONG, W., BUTLER, E., SUN, H., LIU, F., DONG, L., DEVOS, J., GARGIULLO, P., and OTHERS, “Cross-reactive antibody responses to the 2009 pandemic H1N1 influenza virus,” *New England Journal of Medicine*, vol. 361, no. 20, pp. 1945–1952, 2009.
- [21] HEALY, J. D., “A note on multivariate CUSUM procedures,” *Technometrics*, vol. 29, no. 4, pp. 409–412, 1987.

- [22] HUTWAGNER, L. C., THOMPSON, W. W., SEEMAN, G. M., and TREADWELL, T., “A simulation model for assessing aberration detection methods used in public health surveillance for systems with limited baselines,” *Statistics in Medicine*, vol. 24, no. 4, pp. 543–550, 2005.
- [23] INFLATIONDATA.COM, “Historical cpi-u data from 1913 to the present,” http://www.inflationdata.com/Inflation/Consumer_Price_Index/HistoricalCPI.aspx (Accessed November 29, 2009).
- [24] JIANG, W., HAN, S. W., TSUI, K.-L., and WOODALL, W. H., “Spatiotemporal surveillance methods in the presence of spatial correlation,” *Statistics in Medicine*, vol. 30, no. 5, pp. 569–583, 2011.
- [25] KIM, S.-H., ALEXOPOULOS, C., TSUI, K.-L., and WILSON, J. R., “A distribution-free tabular cusum chart for autocorrelated data,” *IIE Transactions*, vol. 39, no. 3, pp. 317–330, 2007.
- [26] KIM, S.-H. and DIEKER, A. B., “Selecting the best by comparing simulated systems in a group of three,” in *Proceedings of the 2011 Winter Simulation Conference* (JAIN, S., CREASEY, R. R., HIMMELSPACH, J., WHITE, K. P., and FU, M., eds.), (Phoenix, AZ, USA), December 2011.
- [27] LAWSON, A. B., *Statistical methods in spatial epidemiology*. John Wiley & Sons Sussex, 2001.
- [28] LEE, M. L., GOLDSMAN, D., KIM, S.-H., and TSUI, K.-L., “Spatiotemporal biosurveillance with spatial clusters: control limit approximation and impact of spatial correlation,” *IIE Transactions*, in print, 2013. DOI: 10.1080/0740817X.2013.785296.

- [29] LONGINI, I. M., NIZAM, A., XU, S., UNGCHUSAK, K., HANSHAOWORAKUL, W., CUMMINGS, D. A. T., and HALLORAN, M. E., “Containing pandemic influenza at the source,” *Science*, vol. 309, no. 5737, pp. 1083–1087, 2005.
- [30] LONGINI JR, I. M., HALLORAN, M. E., NIZAM, A., and YANG, Y., “Containing pandemic influenza with antiviral agents,” *American Journal of Epidemiology*, vol. 159, no. 7, pp. 623–633, 2004.
- [31] LUCAS, J. M. and SACCUCCI, M. S., “Exponentially weighted moving average control schemes: properties and enhancements,” *Technometrics*, pp. 1–12, 1990.
- [32] MEDLOCK, J. and GALVANI, A. P., “Optimizing influenza vaccine distribution,” *Science*, vol. 325, no. 5948, pp. 1705–1708, 2009.
- [33] MELTZER, M., COX, N., and FUKUDA, K., “The economic impact of pandemic influenza in the united states: priorities for intervention.,” *Emerging Infectious Diseases*, vol. 5, no. 5, pp. 659–671, 1999.
- [34] MONTGOMERY, D. C., *Introduction to Statistical Quality Control*. 6th ed., John Wiley & Sons Inc., New York, 2009.
- [35] NEW SCIENTIST, “Warm weather may not halt swine flu,” <http://www.newscientist.com/article/dn17100-warm-weather-may-not-halt-swine-flu.html> (Accessed November 13, 2009).
- [36] NEW YORK TIMES, “City parents opting out of swine flu vaccine,” http://www.nytimes.com/2009/10/29/nyregion/29vaccine.html?_r=1 (Accessed November 6, 2009).
- [37] NEW YORK TIMES, “A nation battling flu and short vaccine supplies,” <http://www.nytimes.com/2009/10/26/health/26flu.html?hpw> (Accessed October 26, 2009).

- [38] NEW YORK TIMES, “Swine flu shots revive a debate about vaccines,” <http://query.nytimes.com/gst/fullpage.html?res=9903EFDA1739F935A25753C1A96F9C8B63> (Accessed October 21, 2010).
- [39] PAGE, E. S., “Continuous inspection schemes,” *Biometrika*, vol. 41, no. 1, pp. 100–115, 1954.
- [40] PAGE, E. S., “Cumulative sum charts,” *Technometrics*, vol. 3, pp. 1–9, 1961.
- [41] PIGNATIELLO, J. J. and RUNGER, G. C., “Comparisons of multivariate CUSUM charts,” *Journal of Quality Technology*, vol. 22, pp. 173–186, 1990.
- [42] PUBLIC HEALTH AGENCY OF CANADA, “H1N1 flu virus in Canada,” http://www.phac-aspc.gc.ca/alert-alerte/h1n1/faq_rg_h1n1-eng.php (Accessed October 26, 2009).
- [43] ROGERSON, P. A., “Surveillance systems for monitoring the development of spatial patterns,” *Statistics in Medicine*, vol. 16, no. 18, pp. 2081–2093, 1997.
- [44] ROGERSON, P. A., “Formulas for the design of CUSUM quality control charts,” *Communications in Statistics — Theory and Methods*, vol. 35, pp. 373–383, 2006.
- [45] ROGERSON, P. A. and YAMADA, I., “Approaches to syndromic surveillance when data consist of small regional counts,” *Morbidity and Mortality Weekly Report*, vol. 53, pp. 79–85, 2003.
- [46] ROGERSON, P. A. and YAMADA, I., “Monitoring change in spatial patterns of disease: comparing univariate and multivariate cumulative sum approaches,” *Statistics in Medicine*, vol. 23, no. 14, pp. 2195–2214, 2004.
- [47] SALKIND, N. J., *Encyclopedia of measurement and statistics*. Thousand Oaks, California, 2007.

- [48] SANDER, B., NIZAM, A., GARRISON JR, L. P., POSTMA, M. J., HALLO-
RAN, M. E., and LONGINI JR, I. M., “Economic evaluation of influenza pan-
demic mitigation strategies in the united states using a stochastic microsimula-
tion transmission model,” *Value in Health*, vol. 12, no. 2, pp. 226–233, 2009.
- [49] SONESSON, C., “A CUSUM framework for detection of space–time disease clus-
ters using scan statistics,” *Statistics in Medicine*, vol. 26, no. 26, pp. 4770–4789,
2007.
- [50] SONESSON, C. and BOCK, D., “A review and discussion of prospective statistical
surveillance in public health,” *Journal of the Royal Statistical Society: Series A*,
vol. 166, no. 1, pp. 5–21, 2003.
- [51] STATISTICS CANADA, “2006 census of population, cansim using chass: 2006
profile of census tracts / age and sex,” <http://dc1.chass.utoronto.ca/cansimdim/English/> (Accessed May 6, 2009).
- [52] STATISTICS CANADA, “2006 census of population, cansim using chass: 2006
profile of census tracts / labour market activity,” <http://dc1.chass.utoronto.ca/cansimdim/English/> (Accessed May 6, 2009).
- [53] STATISTICS CANADA, “2006 census of population, cansim using chass: Labour
force survey estimates, by census metropolitan area, sex and age group, annually
(persons unless specified),” <http://dc1.chass.utoronto.ca/cansimdim/English/>
(Accessed May 6, 2009).
- [54] STATISTICS CANADA, “2006 census of population, cansim using chass: Pro-
file tables/marital status, families & household living arrangements, and
housing,” <http://dc1.chass.utoronto.ca/census/index.html> and <http://www40.statcan.gc.ca/l01/cst01/famil54b-eng.htm> (Accessed May 6, 2009).

- [55] STATISTICS CANADA, “2006 census of population, cansim using chass: Table 510036 -estimates of population, by sex and age group, census divisions and census metropolitan areas, 2006 census boundaries, annually (persons unless specified),” <http://dc1.chass.utoronto.ca/cansimdim/English/> (Accessed May 6, 2009).
- [56] STATISTICS CANADA, “Industry - north american industry classification system 2002 (433), class of worker (6) and sex (3) for the labour force 15 years and over of canada, provinces, territories, census metropolitan areas and census agglomerations, 2006 census - 20 percent sample data,” <http://www.statcan.gc.ca/bsolc/olc-cel/olc-cel?lang=eng&catno=97-559-X2006009> (Accessed Nov 29, 2009).
- [57] STROUD, P., DEL VALLE, S., SYDORIAK, S., RIESE, J., and MNISZEWSKI, S., “Spatial dynamics of pandemic influenza in a massive artificial society,” *Journal of Artificial Societies and Social Simulation*, vol. 10, no. 4, p. e9, 2007.
- [58] SWIVEL PREVIEW, “Consumer price index for medical care,” http://www.swivel.com/data_columns/spreadsheet/1875739 (Accessed November 29, 2009).
- [59] TANGO, T., “A class of tests for detecting general and focused clustering of rare diseases,” *Statistics in Medicine*, vol. 14, no. 21-22, pp. 2323–2334, 1995.
- [60] TSUI, K.-L., CHIU, W., GIERLICH, P., GOLDSMAN, D., LIU, X., and MASCHEK, T., “A review of healthcare, public health, and syndromic surveillance,” *Quality Engineering*, vol. 20, pp. 435–450, 2008.
- [61] TSUI, K.-L., HAN, S. W., JIANG, W., and WOODALL, W. H., “Likelihood ratio methods for spatial and spatiotemporal surveillance,” *IIE Transactions*, p. in press.

- [62] TUIITE, A. R., FISMAN, D. N., KWONG, J. C., and GREER, A. L., “Optimal pandemic influenza vaccine allocation strategies for the canadian population,” *PloS one*, vol. 5, no. 5, p. e10520, 2010.
- [63] TUIITE, A. R., GREER, A. L., WHELAN, M., WINTER, A. L., LEE, B., YAN, P., WU, J., MOGHADAS, S., BUCKERIDGE, D., POURBOHLOUL, B., and FISMAN, D. N., “Estimated epidemiologic parameters and morbidity associated with pandemic H1N1 influenza,” *Canadian Medical Association Journal*, vol. 182, no. 2, pp. 131–136, 2010.
- [64] WOODALL, W. H., “The use of control charts in health-care and public-health surveillance,” *Journal of Quality Technology*, vol. 38, no. 2, pp. 89–104, 2006.
- [65] WOODALL, W. H. and NCUBE, M. M., “Mutivariate CUSUM quality control procedures,” *Technometrics*, vol. 27, pp. 285–292, 1988.
- [66] WOODALL, W. H., SPITZNER, D. J., MONTGOMERY, D. C., and GUPTA, S., “Using control charts to monitor process and product quality profiles,” *Journal of Quality Technology*, vol. 36, no. 3, pp. 309–320, 2004.
- [67] WORLD HEALTH ORGANIZATION, “Pandemic (H1N1) 2009 - update 71,” http://www.who.int/csr/don/2009_10_23/en/index.html (Accessed October 26, 2009).
- [68] WORLD HEALTH ORGANIZATION, “World now at the start of 2009 influenza pandemic,” http://www.who.int/mediacentre/news/statements/2009/h1n1_pandemic_phase6_20090611/en/index.html (Accessed October 26, 2009).
- [69] YANG, Y., SUGIMOTO, J. D., HALLORAN, M. E., BASTA, N. E., CHAO, D. L., MATRAJT, L., POTTER, G., KENAH, E., and LONGINI, I. M., “The transmissibility and control of pandemic influenza A (H1N1) virus,” *Science*, vol. 326, no. 5953, pp. 729–733, 2009.

- [70] ZHU, Y., WANG, W., ATRUBIN, D., and WU, Y., “Initial evaluation of the early aberration reporting system—florida,” *Morbidity and Mortality Weekly Report*, vol. 54, pp. 123–130, 2005.