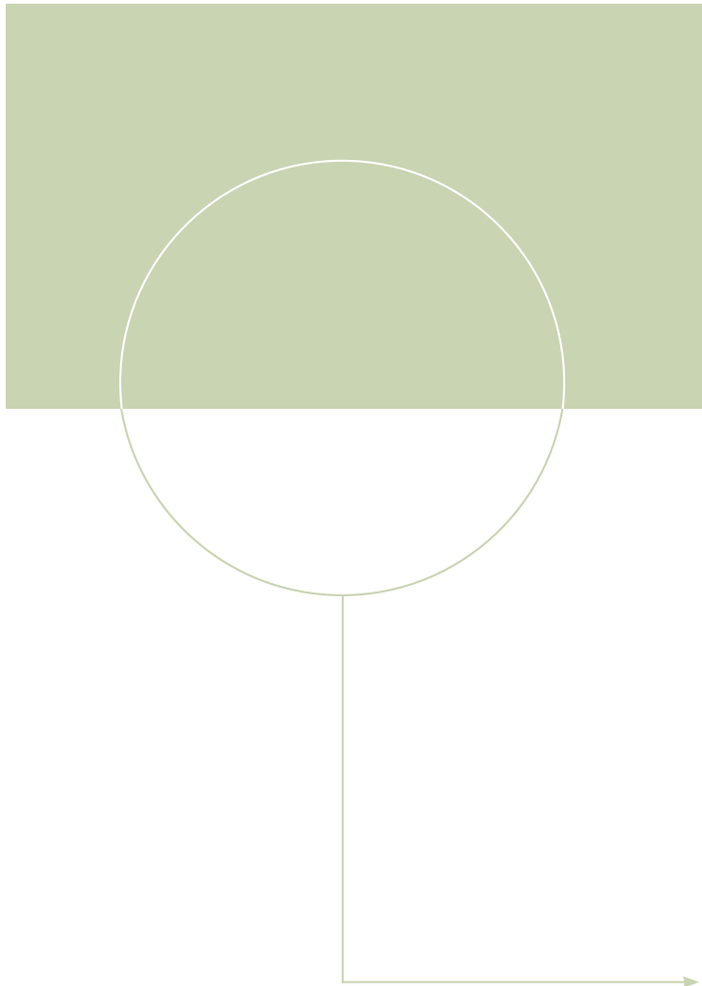

Master's thesis

2020

Martin Willoch Olstad

NTNU
Norwegian University of
Science and Technology
Faculty of Economics and Management
Department of Industrial Economics and Technology
Management

Martin Willoch Olstad

Optimizing Resource Allocation during Epidemic Outbreaks: An Approximate Dynamic Programming Approach for Cholera

July 2020

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of
Science and Technology

NTNU

Norwegian University of
Science and Technology

Optimizing Resource Allocation during Epidemic Outbreaks: An Approximate Dynamic Programming Approach for Cholera

Martin Willoch Olstad

Industrial Economics and Technology Management
Submission date: July 2020
Supervisor: Ruud Egging-Bratseth
Co-supervisor: Henrik Andersson

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Preface

This Master's Thesis concludes my Master of Science in Industrial Economics and Technology Management with a specialization in Managerial Economics and Operations Research at the Norwegian University of Science and Technology (NTNU).

Throughout the last semester, epidemiology received significant attention as the COVID-19 pandemic devastated countries all over the world. Most of this thesis was written in volunteered quarantine, watching healthcare systems struggling with surging demand. Now, possibly more than ever, an efficient allocation of scarce medical resources is vital. The potential impact of this research field has been a genuine motivational factor when writing this thesis. There are no known effective antiviral medications or vaccines for COVID-19 as of today. However, cholera is both treatable and preventable. Therefore, I hope this grave situation we have experienced during the past months, increase awareness of other infectious diseases and mobilize funds and attention to eradicating diseases we have the knowledge to defeat together.

I would like to express my sincere gratitude to my supervisors, Professor Ruud Egging-Bratseth and Professor Henrik Andersson, for contributing with inspiring discussions, valuable perspectives and constructive feedback. I would also like to extend my gratitude to Dr. Jong-Hoon Kim and Dr. Vittal Mogasale at the International Vaccine Institute for providing insightful feedback on our epidemiological results. Lastly, I would like to thank family and friends for support, creative ideas and endless proofreading.

Oslo, July 2, 2020.

Martin Willoch Olstad

Abstract

Epidemic outbreaks affect the lives of people all around the world. The COVID-19 pandemic painfully demonstrates that despite the major medical achievements in the past centuries, there is still a need for efficient responses to epidemics. Even for diseases where there exists both effective vaccines and medications, such as cholera, epidemic outbreaks occur and costs many lives every year. Efficient responses with a functioning strategy are vital when dealing with epidemics, and this thesis aims to provide decision-support both in advance of and during epidemic outbreaks on what response policy might be most effective.

This thesis proposes a resource allocation model combined with a cholera epidemic model. Together, the models aim to allocate medical intervention resources to save as many lives as possible during cholera outbreaks, within the constraints of available medical personnel and temporary medical facilities. The proposed resource allocation model is solved using an approximate dynamic programming (ADP) approach with a neural network as an approximation technique for the value function. The resource allocation problem at the start of each time period is in turn solved heuristically using a local search procedure. The cholera model combines and extends previous works to a multi-region, multi-intervention Susceptible-Asymptomatic-Infected-Recovered-Bacteria (SAIR-B) model. Environmental fluctuations are an important factor in spreading cholera, and it has been linked to climatic conditions. Therefore, the bacteria dispersal rate between regions in the SAIR-B model is included as a stochastic variable in the ADP model.

The epidemic model is calibrated to the 2010 cholera outbreak in Haiti, and a computational study is conducted on the calibrated epidemic model and on an alternative epidemic model with higher, but still realistic, bacteria excretion rate. The results indicate that the value function approximation converges towards a consistent ADP policy, and that this policy is robust to various bacteria dispersal rate distributions. The investigation regarding the availability of medical resources indicates that although rehydration solution is essential to treat symptomatic cholera-infections, additional vaccines made a larger impact on the total disease-induced fatalities. Increased availability of vaccines also appeared to reduce the fatalities more than earlier arrival of vaccines. Decision-makers should thus be aware that focusing on collecting reliable surveillance data to get a sufficient overview of the entire outbreak situation, may prove more important than requesting vaccines as rapidly as possible from the International Coordinating Group on Vaccine Provision.

Sammendrag

Epidemiutbrudd påvirker mennesker over hele kloden. COVID-19 pandemien viser at til tross for store mediniske framskritt de siste tiårene, er det stadig behov for en effektiv respons ved epidemiutbrudd. Selv for sykdommer som det finnes fungerende vaksiner og medikamenter mot, som kolera, forekommer epidemiutbrudd som koster liv hvert eneste år. Effektiv respons med en velfungerende strategi er essensielt når man skal respondere på en epidemi. Målet ved denne oppgaven er å utvikle modeller for beslutningsstøtte til valg av strategier som gir en mest mulig effektiv respons. Støtten kan gis både i forkant av og under et epidemiutbrudd.

Oppgaven foreslår en ressursallokerings-modell kombinert med en kolera-modell. Modellene allokerer medisinske intervensjonsressurser for å redde så mange liv som mulig under kolerautbrudd. Ressursene allokeres innenfor begrensningene gitt av tilgjengelig medisinsk personnel og midlertidige medisinske fasiliteter. Den foreslåtte ressursallokeringsmodellen er løst med en approksimert dynamisk programmerings-tilnærming, med et nevralt nettverk som approksimeringsteknikk for verdifunksjonen. Ressursallokeringsproblemet i begynnelsen av hver tidsperiode løses heuristisk med en lokal søkprosedyre. Kolera-modellen kombinerer og videreutvikler modeller fra eksisterende forskning til en multi-region og multi-intervensjons Susceptible-Asymptomatic-Infected-Recovered-Bacteria (SAIR-B) modell. Variasjoner i lokalt klima og miljø er en sentral faktor ved spredning av kolera. Derfor er bakteriespredningsraten mellom regioner i SAIR-B-modellen inkludert som en stokastisk variabel i ADP-modellen.

Den epidemiologiske modellen er kalibrert etter kolerautbruddet på Haiti i 2010, og et numerisk studie er gjennomført på både den kalibrerte modellen og en alternativ epidemimodell med høyere, men fortsatt realistisk, ekskresjonsrate av bakterier. Resultatet indikerer at den approksimerte verdifunksjonen konvergerer mot en konsekvent ADP beslutningsregel, og at beslutningsregelen er robust for variasjoner i ekskresjonsrate-distribusjonen. Undersøkelsen av tilgjengelighet av medisinske ressurser indikerer at selv om rehydreringsløsninger er essensielt når man behandler kolera-pasienter med symptomer, så reduserte en økning av vaksiner dødeligheten mer. Det viste seg også at en økning i tilgjengelighet av vaksiner betydde mer for reduksjon av antall døde enn raskere tilgang på vaksiner gjorde. Dette indikerer at beslutningstakere burde fokusere på å samle inn pålitelige data om epidemispredningen for å få tilstrekkelig oversikt over utbruddet, framfor å forespørre vaksiner så raskt som mulig fra International Coordinating Group on Vaccine Provision.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Throughout human history, there has been a wide range of severe outbreaks of infectious diseases, almost eradicating entire populations. As the world became more interconnected through trading routes, these diseases spread even more quickly. During the cholera epidemic in London in 1854, John Snow became the modern-day father of epidemiology, when he connected the outbreak to a specific water pump on Broad Street (Kanchanaraksa, 2008). Since then, there have been developed quantitative models aiming to explain and predict the distribution and determinants of diseases among populations.

Today, the field of medicine and epidemiology ensures efficient response to epidemics in many countries. Still, infectious diseases are spreading faster and further than ever before (World Health Organization, 2018, pp. 17). During the past 20 years, we have faced cholera outbreaks in Haiti and Yemen, Ebola outbreaks in the western parts of Africa, the influenza pandemic swine flu in 2009 and the ongoing COVID-19 pandemic. While the new outbreaks can be discouraging, there have been significant humanitarian achievements in the past decades. In 1979 the World Health Assembly declared that the devastating smallpox disease was eradicated (World Health Organization, 2019d) and since the polio eradication initiative started in 1988, cases due to wild poliovirus have decreased by more than 99% (World Health Organization, 2019b).

One reason for the eradication of diseases is the development of efficient vaccines. Even with novel pathogens, the global medical community is quick to develop new vaccines, thanks to the collective effort. However, even as vaccines are available in large proportions of the world, diseases that could be vaccinated against and properly treated for, keep devastating developing countries. An example of this is cholera, where, despite having

both oral vaccines and proper medical treatment against it, the number of cases has remained high the past decades, as seen in Figure 1.1. There are still explosive outbreaks in regions with inadequate access to clean water and sanitation. While the progress in medicine ensures proper vaccination and treatment of the population in industrial countries, it does not help the most vulnerable populations of financially weak developing countries.

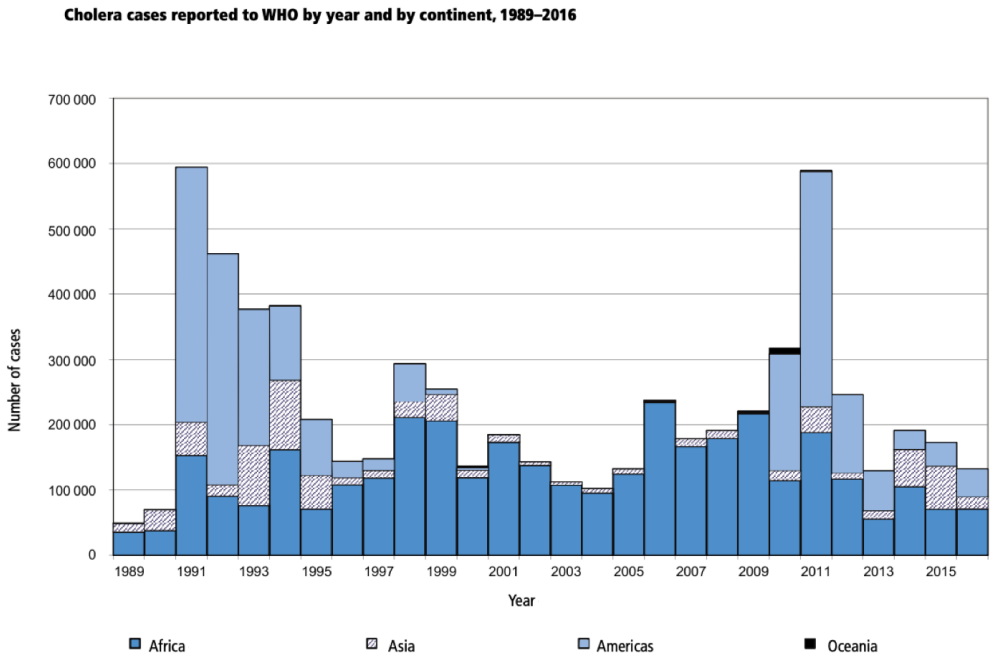

Figure 1.1: Officially reported cholera cases the past decades, modified from World Health Organization (2017c).

The substantial amount of research on epidemiological modeling has contributed to significant improvements in epidemic control. Separately, within the field of operations research, there has been considerable research on resource allocation during emergencies and disasters. However, there has been limited work on combining these approaches to a single decision-support tool for epidemic response. The goal of this thesis is to develop an integrated epidemic and resource allocation model that efficiently allocates medical resources during an infectious disease outbreak. The focus is on diseases with known and well-researched prevention and treatment methods within financially weak regions. A case study is conducted on cholera outbreaks in Haiti. The purpose of the model is to help governments and non-government organizations (NGO) make informed decisions to ensure an immediate, efficient response to an outbreak.

This thesis begins with an introduction to the different phases of an epidemic and the transmission dynamics of cholera, in Chapter 2. In Chapter 3, a literature review is conducted, focusing on previous emergency and epidemic logistics literature and epidemiological modeling. Chapter 4 describes the relevant theory to understand the epidemic model and resource allocation model presented in this thesis. The dynamic epidemic response resource allocation problem is described in Chapter 5. Chapter 6 presents mathematical models aiming to solve this problem. In Chapter 7, different solution methods for solving the problem are proposed. Chapter 8 reports the data used during the analysis of the cholera outbreaks in Haiti. Chapter 9 presents and discusses the computational results from applying the cholera outbreak case on the developed models. Finally, Chapter 10 summarizes the findings of this thesis, identifies the weaknesses of the modeling approach and proposes further research to better address these challenges in the future.

Chapter 2

Background

This chapter describes relevant epidemiological and logistical aspects for response in the case of an emergency, and epidemics in specific. First, Section 2.1 describes different characteristics of emergencies and various diseases. The different phases of an emergency and the decisions typically made during each phase from the perspective of a decision-maker, such as local government and NGOs, are described in Section 2.2. Then, Section 2.3 describes the relevant transmission dynamics of a cholera epidemic. Lastly, Section 2.4 describes different intervention methods to respond and contain cholera epidemics.

2.1 Characteristics of Emergencies and Diseases

In this section, characteristics for various emergencies and infectious diseases are discussed. Emergencies with varying characteristics require different response measures. This should be reflected in the type of modeling approach taken when conducting research. Thus, to understand what aspects of the current emergency logistics literature is relevant, it is essential to understand the differences among distinct emergencies.

2.1.1 Emergency characteristics

According to the Cambridge Dictionary, an *emergency* is *"a dangerous or serious situation, such as an accident, that happens suddenly or unexpectedly and needs immediate action"* (Cambridge Academic Content Dictionary, 2009). This definition stresses the urgency of the situation. However, it does not clarify the magnitude of the event. Throughout this thesis, the word *emergency* is used when referring to big-impact events, affecting

entire communities, and not individual emergencies, such as cardiac arrests and strokes.

All large-scale emergencies share certain characteristics. They happen abruptly and results in a sudden surge in demand for relief supplies and services. However, different types of emergencies and disasters require different preparedness and response measures.

Certain emergencies can, to some extent, be anticipated and thus be better prepared for. Hurricanes in the Atlantic Basin occur seasonally from June to November (National Hurricane Center, 2020) and certain kinds of diseases also have seasonal attributes. Cholera outbreaks in Bangladesh have, for instance, been connected to the monsoon season (Emch et al., 2008). Seasonality of emergencies can support the preparation, but even though seasonal events happen within a certain time-frame, the exact location of where the emergency occurs can be difficult to anticipate. For other kinds of emergencies, it is the opposite. Large-scale volcano eruptions and earthquakes occur relatively infrequently; however, due to geological constraints, the possible locations of such emergencies are easier to anticipate. Volcano eruptions occur where there are active volcanoes, and earthquakes usually, though not always, occur along tectonic plate interaction zones. Certain emergencies are both temporally and spatially difficult to anticipate, such as terrorist attacks.

The progression of an emergency is also a varying characteristic. An earthquake can typically be felt for a few seconds (GNS Science, 2020), although aftershocks may occur. The earthquake may cause massive damage in several regions outside the epicenter, but it does not evolve over extended periods of time. A hurricane strikes and moves through several different regions. While hurricanes can be properly prepared for, there is not possible to contain hurricanes and they will live through their life cycles independent of human intervention. Other emergencies, such as epidemic and wildfire outbreaks, evolves stochastically and can last for significantly longer than earthquakes and individual hurricanes. Epidemics and wildfires should be contained through human intervention, if not, they can expand into drastically larger areas, causing significantly more damage. Thus, while most emergencies require relief distribution, certain emergencies require an immediate and efficient response to contain the emergencies and limit their extent.

2.1.2 Disease characteristics

The term *disease* may refer to a substantial amount of different phenomena and is defined by the Cambridge Dictionary as *"an illness caused by an infection or by a failure of health and not by an accident"* (Cambridge Learner's Dictionary, 2007). Thus, it may refer to non-communicable diseases, that is, diseases not transmissible directly between individuals, such as Alzheimer's disease. The term *disease* can also refer to communicable diseases, also called infectious diseases, meaning they can be transmitted between individuals. This thesis focus on response to diseases transmissible among people, and

thus communicable diseases are of interest. If a communicable disease rapidly spreads to a large number of individuals, the event is known as an *epidemic*. Large epidemics can also evolve into *pandemics*, spreading into multiple large regions and continents.

Different diseases have different methods of transmission. Certain diseases transmit by contact through air droplets from the respiratory system of the infected individual, with influenza being a familiar example. Fecal-oral transmission is caused by the ingestion of fecal material from an infected individual, for instance, drinking fecally contaminated water. Examples of fecal-oral transmittable diseases include cholera and polio. Vector-borne diseases are caused by pathogens living inside organisms other than humans, such as insects, and typically transmits to humans through insect bites. Malaria and dengue fever are both examples of vector-borne diseases transmitted through mosquito bites. In addition, diseases such as HIV can be transmitted sexually, by blood and vertically, meaning carried on from mother to child (Checchi, 2009).

Various pathogenic microorganisms, for instance, viruses, bacteria and parasites, may cause infectious diseases (World Health Organization, 2016). The most efficient response depends on the kind of pathogenic microorganism. For instance, antibiotics are antibacterial, and thus ineffective against viral infections. The symptoms of the specific disease are also important when treating it. If the disease cause dehydration, then rehydration treatment is necessary. In addition, supportive treatment, such as painkillers, can be applied to increase the quality of life for infected and symptomatic individuals.

Vaccines are substances developed to increase disease immunization by helping the human body's immune system to recognize and fight pathogens (World Health Organization, 2019c). There are different methods of injecting vaccines. Intramuscular and subcutaneous vaccines are injected using a syringe, scarification vaccines are injected through a skin scratch and oral vaccines are mixed with drinking water and orally ingested. While all forms of vaccination require proper dosage, the latter does not require medical personnel to properly inject the vaccine.

For certain diseases, efficient vaccines and specific medical treatment are not available. While there is extensive research on developing an efficient malaria vaccine, there is yet to be a commercially available vaccine (World Health Organization, 2017a), and although there is an Ebola vaccine, there is no antiviral drug available in case of infection (Centers for Disease Control and Prevention, 2019).

Diseases with both commercially available vaccines and treatment may still pose a threat. Examples include pandemic influenza and cholera. These diseases may still pose a threat due to lack of medication and vaccination production capacity, basic sanitation infrastructure and lack of market access, respectively (World Health Organization, 2018).

This thesis focuses on fecal-oral transmittable diseases, with some degree of anticipa-

tion of outbreaks due to seasonality or poor infrastructure, and existing treatment and vaccines, though in scarce amounts. These factors affect both the epidemiological and mathematical programming model presented in this thesis. However, through minor adjustments and extensions, the models may be used for other communicable diseases.

2.2 Epidemic Phases

The WHO describes the typical epidemic phases as introduction or emergence, localized transmission, amplification and reduced transmission (World Health Organization, 2018, pp. 28), as is shown in Figure 2.1. First, the disease is introduced to a community. During the second phase, there are sporadic infections within the community. In the third phase, the infections amplify and it turns into an epidemic. During this phase, the reproduction number is high, that is, the expected number of new cases of infected generated by one infected individual. Throughout the epidemic, the population develops increased immunity due to effective interventions and recoveries from the disease, the reproduction number falls, constituting the fourth and final phase, reduced transmission immunity. When the reproduction number is consistently below 1.0, meaning each infected person on average infects less than one other person, the epidemic fades out and ends.

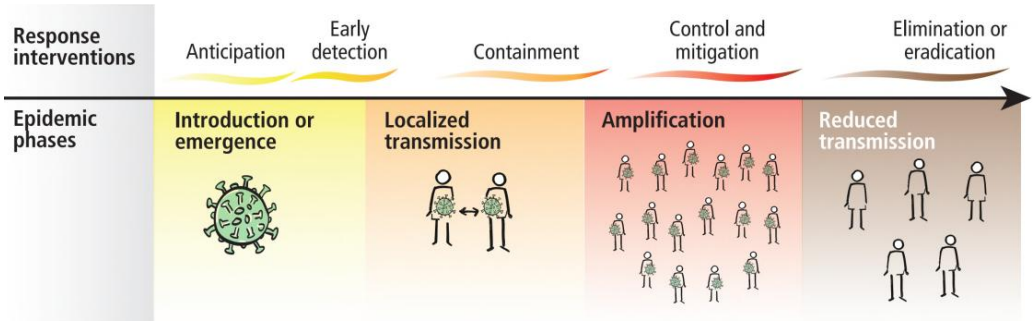

Figure 2.1: Typical phases during an epidemic outbreak, reproduced from World Health Organization (2018).

To aid decision-makers and responders to an epidemic during the four phases, the WHO has developed a general framework for epidemic responses that can be applied to several different diseases. The *Coordinating responders - Health Information - Communicating risk - Health Interventions* framework functions as a checklist to avoid overlooking essential aspects during an epidemic outbreak (World Health Organization, 2018, pp. 31).

Coordinating responders: During an emergency, there are often several organizations involved in responding. To ensure the best possible use of the available resources, a well-coordinated response among the responders is essential. To achieve such effective

coordination, the WHO proposes establishing a common emergency operation center, developing a joint plan of action that is regularly updated and includes the distribution of roles and responsibilities among the responders and tools to ensure communication among the responders, such as contact information directories.

Health Information: The WHO defines two key types of information during an epidemic: surveillance of the disease and information on the interventions. In addition to collecting the data, the framework stresses the importance of a common understanding and definition is essential among the different responders. Questions that should be answered include if the case definition is shared among all stakeholders, what the risk groups are, what resources, both material and human, are available and how much is required, and what are indicators of success.

Communicating risk: In addition to the actual epidemic, a new challenge has arisen with the use of social media. The rapid spread of information cause what the WHO has termed an *infodemic*: a rapid spread of both reliable and unreliable information in parallel to the rapid spread of the actual disease. To ensure that citizens listen to the governments and take the necessary precautions to avoid any escalation of the epidemic, proper risk communication is essential. WHO stresses that this requires two-way communication, quickly communicating protective measures that people can take through mass media, but also listening to the concerns and perceptions of the population.

Health Interventions: Different diseases have different characteristics, and thus require different interventions to ensure containment. The health intervention should ensure that critical interventions needed to control the outbreak are both mapped and adequately implemented, and how the interventions impact the epidemic spreading dynamics.

2.3 Cholera Transmission

Cholera is an acute diarrhoeal disease caused by the bacterium *Vibrio cholerae* and individuals are infected by ingesting food or water contaminated by the bacteria (World Health Organization, 2019a). Once infected, the acute watery diarrhea causes dehydration and can be fatal within hours if the infected individual is not treated.

The incubation period, that is, the period between the individual is exposed to the pathogen to the first symptoms are showing, is between 12 hours and 5 days. However, about 80% of the individuals infected with cholera are asymptomatic, i.e. do not develop symptoms. The bacteria are present in their body for up to 14 days and brought back to the environment through their feces (World Health Organization, 2018, pp. 165). Without proper sanitation, the feces can contaminate the local water-source and the immediate surroundings, causing more infections. The infected individuals that do develop symptoms can experience mild, moderate and severe dehydration (World Health Organization, 2019a).

After recovery, a limited, natural immunity to the bacteria is developed, lasting from 6 months to several years, depending on the response of the individual's immune system.

Today, cholera can be found around the globe in both endemic and epidemic states. The WHO defines a cholera-endemic if the area has confirmed cholera cases the past 3 years with evidence of local transmission, that is, it was not brought to the area through migration from elsewhere (World Health Organization, 2019a). Endemic areas can also experience epidemic outbreaks, defined by the number of cases being higher than expected.

Cholera outbreaks can be both sporadic and seasonal. Outbreaks can occur as the result of migration, but also through the water-source network. Given an outbreak in an area, the water-source can have an increase in the concentration of cholera bacteria, which is brought to other areas by rivers. The cholera bacterium can persist for long periods of time in an aquatic environment. In addition, the bacteria can survive on fish, shellfish and zooplankton (Colwell, 1996). If the fish is later eaten raw, it may cause cholera infections that can develop into outbreaks. The cholera bacteria may also multiply and persist in moist food for a long period of time (World Health Organization, 2018, pp. 163). The seasonal cholera outbreaks can occur both in dry seasons and rainy seasons. During dry seasons, the absence of many water-sources causes a single contaminated water-source to infect a large number of people. While during rainy seasons, the rainfall can disperse the contaminated feces into multiple water-sources (Olson et al., 2018, pp. 12).

While the cholera outbreaks can be tied to water-sources, human migration can also be a significant factor. Chin et al. (2011) investigated the origin of the Haitian cholera outbreak in October 2010 and concluded that cholera was introduced to Haiti as a result of human activity, likely from a distant geographic source.

The typical duration of a cholera outbreak depends on the population density and the population number in the area. In urban settings, the duration can be between 2 to 4 months, reaching the peak number of infected after 2 to 8 weeks, while the duration in rural settings typically is 3 to 6 months with the peak reached after 1 to 3 months. In refugee camps, the typical epidemic duration is 1 to 3 months, with the peak reached after 2 to 4 weeks (Olson et al., 2018, pp. 12).

2.4 Cholera Intervention

Cholera is often described as an important indicator of inequity and a lack of social development (World Health Organization, 2019a). This is because there exist both effective vaccines and medication to prevent and properly treat cases of cholera. Yet, there are still outbreaks occurring in the most vulnerable regions, lacking basic sanitary infrastructure.

This section describes the properties and typical applications of key intervention methods in response to cholera outbreaks.

2.4.1 Sanitation infrastructure

The Global Task Force on Cholera Control (GTFCC) published a strategy to eradicate cholera within 2030 and identified that while an emergency response to outbreaks reduces the mortality and morbidity of an outbreak, it does not provide long-term prevention of cholera (Global Task Force on Cholera Control, 2017, pp. 7). To prevent and eradicate the disease, the development of basic water, sanitation and hygiene services (WASH) and mass-vaccination with the oral cholera vaccine (OCV) is necessary. In the strategy, the GTFCC declared that even though there exist measures to prevent and control cholera, these measures are not used optimally in local contexts and are not supported with sufficient financial and human resources, arguing that more than 80% of cholera-affected countries have reported insufficient financing to meet WASH targets (Global Task Force on Cholera Control, 2017, pp. 8).

Development of sanitation infrastructure is therefore essential to prevent cholera outbreaks. However, being a long-term prevention strategy, not a short-term response to an outbreak, sanitation infrastructure is not the focus of this thesis.

2.4.2 Outbreak alert

The combination of short incubation periods and a very short time until cholera can prove fatal, cause the immediate response to an outbreak to be essential. The short incubation period also causes cholera outbreaks to be particularly explosive, even if most individuals exposed to the bacteria do not develop symptoms. The case-fatality rate (CFR), that is, the ratio between the number of fatalities and cases, in untreated cases may reach 30-50%, while the Global Task Force on Cholera Control (GTFCC) stated the cholera CFR benchmark to be below 1% (Global Task Force on Cholera Control, 2004, pp. 7).

While there can be several reasons for diarrhea in an area, if the number of cases is especially high, a cholera outbreak should be suspected. Immediately after, the preparation of an on-site investigation should be initiated. The investigation should be undertaken within 24 hours of the alert and consists of taking samples of individuals with acute diarrhea (Olson et al., 2018, pp. 19). It takes around two days to get results from laboratories. However, the GTFCC advises not to delay the treatment of individuals showing symptoms of cholera after the laboratory has confirmed the outbreak. There are no available rapid tests to measure the concentration of cholera bacteria in water. However, the general quality of water sources can be measured during an on-site investigation with results within minutes by measuring the pH, turbidity and concentration of free residual

chlorine (Olson et al., 2018, pp. 36). An overview of the initial outbreak timeline with investigation and vaccine requesting is shown in Figure 2.2.

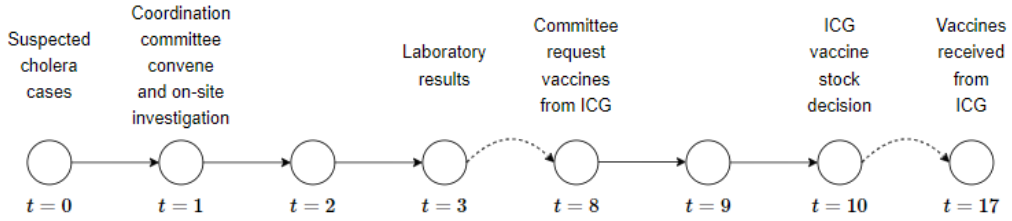

Figure 2.2: Timeline from cholera suspicion to available vaccines. The vaccine requests are elaborated in Section 2.4.4.

2.4.3 Cholera coordination committee

In countries where cholera is relatively common, a *cholera coordination committee* should be appointed and ensure proper preparedness and response to outbreaks. The committee works as the central decision-maker during the outbreak and should facilitate collaboration between the different organizations responding to the outbreak. The GTFCC advises that the committee meets at least once per week during an outbreak. Once an outbreak is suspected, the committee should convene and initiate the immediate response to control a cholera outbreak. The committee should make an inventory of available medical supplies, setting up temporary treatment centers where needed, implementing measures to control the spread, training medical personnel, collect and analyze the data on cases and deaths, and inform and educate the public.

2.4.4 Oral cholera vaccines

There are three recognized OCVs (Global Task Force on Cholera Control, 2017, pp. 8). The OCVs can be applied as a long-term measure to eradicate cholera through mass-vaccination, but a reactive vaccination strategy can also be implemented. Havumaki et al. (2019) showed that a reactive vaccination strategy can prove efficient in controlling cholera outbreaks in crowded areas, such as refugee camps. The OCV takes effect immediately and provides protection against cholera for two to three years, thus working as both an immediate response and long-term prevention measure (Global Task Force on Cholera Control, 2017, pp. 10). In contrast to most vaccines, OCVs does not require intramuscular injection and given av appropriate dosage, it can be taken orally with clean water without the need of medical personnel. Because the OCVs are ingested orally, the distribution of the vaccine is easier, thus differentiating the response strategy from other communicable diseases. However, the OCVs also require refigration, thus they cannot be

distributed to households long before it must be taken.

For Shanchol and Dukoral, two of the three available OCVs, there have been established global stockpiles (Global Task Force on Cholera Control, 2017, pp. 8). The International Coordination Group on Vaccine Provision (ICG) manages the stockpile, which is established for outbreak response purposes, both for reactive vaccination in cholera outbreak areas and pre-emptive vaccination in areas with an increased risk of cholera outbreaks (World Health Organization, 2013, pp. 17). ICG and its partners determine the deployment of OCVs to cholera outbreak areas based on the severity of the outbreak, the potential impact of vaccination and the local capacity to organize a vaccination campaign. Even with the OCV stockpile, the GTFCC identified the insufficient availability of vaccines as an important challenge in the cholera eradication strategy. In 2017, the OCV production capacity was 25 million doses per year, but Vaccine Alliance estimated a global demand of 76 million doses in 2020 (Global Task Force on Cholera Control, 2017, pp. 19).

Olson et al. (2018) estimate one week of preparation for the coordination committee to request vaccines from the ICG. After receiving a vaccine request, the ICG reviews it within two days, and if granted, the vaccine transportation takes around seven days. If a double-dose vaccination strategy is used, which provides long-term immunity, two doses must be taken with two weeks apart. Thus, the time from requesting vaccines to providing lasting immunity for parts of a population may take several weeks. It is possible to shorten the time and increase the coverage at the expense of long-term immunity by administering single doses. This vaccination strategy provides short-term immunization and can thus be efficient against the immediate outbreak.

If the vaccine supply is limited, the committee has to assess whether to employ a reactive or pre-emptive vaccination strategy. That is, whether to provide vaccines to the population currently affected by cholera or the population where the risk of new outbreaks is the highest. The choice of strategy should be based on where the risk of cholera mortality is the highest, what phase the outbreak is currently in and the availability of other intervention methods (Olson et al., 2018, pp. 61).

2.4.5 Oral rehydration solutions

Independent of symptoms, individuals testing positive for cholera should be treated with oral rehydration solutions (ORS), consisting of sugars and salts dissolved in clean water. Without symptoms, an infected individual can ingest the ORS at home without surveillance.

For mild and moderate dehydration, ORS and Ringer lactate (RL) is used under close surveillance. Up to 6 liters of ORS is required for adults during the first day of the rehy-

dration treatment (World Health Organization, 2019a). If a person is showing symptoms of severe dehydration, ORS is used in addition to intravenous (IV) fluid therapy, RL and antibiotics.

2.4.6 Antibiotics

Antibiotics can shorten the duration of the disease (Rahaman et al., 1976), but are only used in the most severe cases due to increasing antimicrobial resistance (Global Task Force on Cholera Control, 2004, pp. 29). According to Andrews and Basu (2011), antibiotics may also decrease the rate of cholera bacteria excretion of infected individuals.

2.4.7 Clean water and disinfectant

If the local water source was the initial cause of the cholera outbreak, it is important to disinfect it, either centrally or at the household level. Even if the water-source was initially not contaminated, it might become so during the progression of an outbreak, due to contaminated articles such as buckets and due to improper disinfectant of hands and bodies of people, even those not showing symptoms, but still exposed to the cholera bacteria. The water can be disinfected using a chlorine solution. For households, boiling the water and disinfecting it using UV lamps are also possible methods. In addition to disinfecting the water-source, the GTFCC advises a distribution of 20 liters of clean water per person per day during a cholera outbreak.

2.4.8 Cholera treatment facilities

During a cholera outbreak, there are three health care facilities specifically for cholera treatment: cholera treatment centers (CTC), cholera treatment units (CTU) and oral rehydration points (ORP) (Olson et al., 2018, pp. 34). The location of the cholera treatment facilities is not necessarily static, and they can be redeployed as the epidemic evolves and affects new areas. The cholera treatment facilities should have a prepositioned medical supply buffer stock sufficient to treat the expected demand for around two weeks.

The CTC is the health facility with the largest treatment capacity, with around 50 to 200 beds, and requires constant staffing. The staffing requirements are 75 medical personnel for 100 beds (MSPP and CDC, 2011, pp. 33). A CTC works as a treatment facility for the most severe cases, where individuals with severe symptoms are relocated to the CTC from more decentralized facilities, but also as a local treatment facility caring for infected individuals showing only mild symptoms. A CTC should be established within or next to an existing health care facility if available If not, a large community building can be transformed into a CTC. The purpose of the CTC is to allow access to the largest possible number of patients, and they should thus be centrally located. A CTC should always be

accessible by road to avoid any mobility issues with ambulances and supply deliveries. A CTC is split into a contaminated zone and a neutral zone. In the contaminated zone, the cholera bacteria are expected to be present in large amounts, while the neutral zone is reserved for staff and supplies. The two zones must be clearly, physically separated. In addition to the separation, the flow of staff and patients should be strict to reduce the risk of contamination of the neutral zone. Olson et al. (2018) propose a CTC separation and flow shown in Figure 2.3.

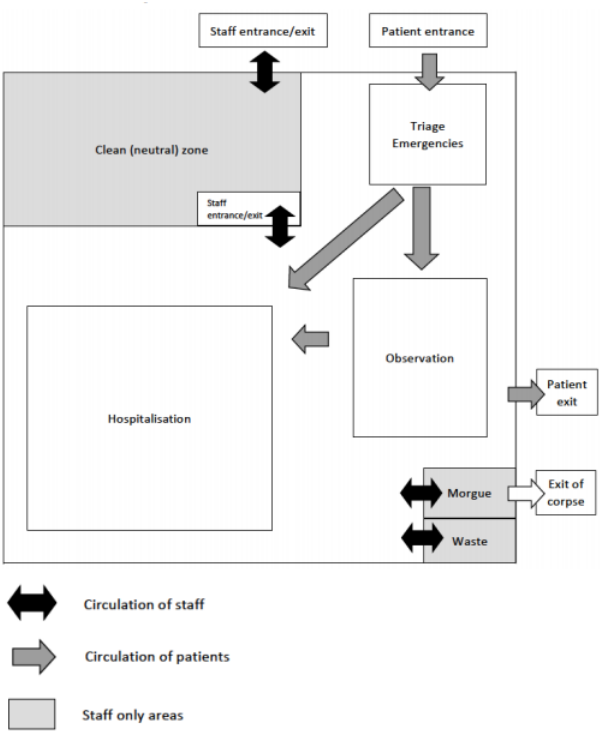

Figure 2.3: Cholera treatment center zones, and staff and patient flows, reproduced from Olson et al. (2018).

The CTU is a smaller health facility, but have the same location requirements as a CTC and can also treat patients requiring IV, in addition to ORS treatment. The capacity is typically 30 beds and CTUs also require constant staffing, with a total of 16 medical personnel (MSPP and CDC, 2011, pp. 33).

The ORP is a small, decentralized facility that distributes ORS to the public and refers to severe symptoms to CTUs or CTCs. An ORP needs only staffing 8 to 12 hours per day. There are typically 3-4 kilometers between every ORP. They can be placed next to existing health care facilities, but also easily accessible by the affected population, for

instance, next to the main road. Both tents and empty buildings can be used and an ORP can become operative within hours after the location is selected (Olson et al., 2018, pp. 115).

Chapter 3

Literature Review

This chapter reviews the literature on epidemic response, control and logistics. Section 3.1 reviews epidemiological modeling, focusing on modeling cholera outbreaks. Section 3.2 provides an overview of emergency logistics within the operations research (OR) field. Section 3.3 reviews emergency logistics literature focusing on epidemic emergencies. The individual subsections are summarized and gaps in the literature are identified in Section 3.4.

3.1 Epidemiological Modeling

The field of epidemiological modeling is almost a century old, beginning when Kermack and McKendrick (1927) introduced the Susceptible-Infected-Recovered (SIR) model, a system of differential equations. Each differential equation represents a compartment, that is, a homogeneous group of the overall population with similar traits. All individuals inside a compartment have the same transition rates to other compartments. Figure 3.1 depicts a simple SIR-model, where individuals in the susceptible compartment transition to the infected compartment with a given infection rate, and the infected individuals recover and transition to the recovered compartment with a given recovery rate.

Since then, more complex models have been developed by including disease-specific mechanisms, yet still relying on similar assumptions as in Kermack and McKendrick (1927), such as a homogeneous population. For cholera, the environment is an particularly important source of transmission. Capasso and Paveri-Fontana (1979) investigated the relationship between the infected individuals and the bacteria concentration in the aquatic reservoir. Codeço (2001) extended the model and introduced the Susceptible-

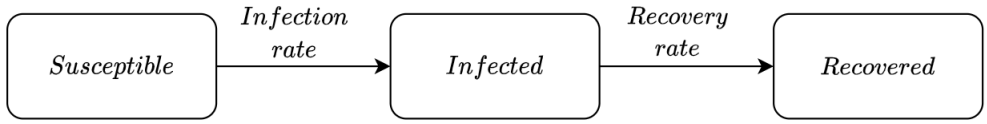

Figure 3.1: Schematic diagram of the Susceptible-Infected-Recovered (SIR) compartmental model.

Infected-Bacteria (SIB) model, which incorporated the environmental transmission factor into the SIR model. By including a compartment for the cholera bacteria concentration in the aquatic reservoirs, the transmission rates, i.e. the rate at which the susceptible population get infected, becomes explicitly dependent on the bacteria concentration in the aquatic reservoirs. Thus, when a susceptible individual becomes infected, it contributes to an increase in cholera bacteria in the reservoirs, which in turn will increase the rates at which other susceptible individuals become infected.

Hyperinfectious bacteria state is another cholera-specific mechanism that can be incorporated into the epidemiological models. Hartley et al. (2005) extend the SIB model by dividing the bacteria concentration compartment into a hyperinfectious and a non-hyperinfectious state. It is assumed that hyperinfectious bacteria are excreted from human individuals, and that these decay and transform into a non-hyperinfectious state.

A large proportion of cholera-infected individuals develop no or mild symptoms. To this end, certain models, such as the ones developed in Neilan et al. (2010) and Andrews and Basu (2011), introduce an additional compartment to differentiate between symptomatic and asymptomatic infections.

The models presented thus far assume homogeneous populations, that is, every individual in a population have the same chance of transitioning. Agent-based modeling (ABM) relax this assumption. Crooks and Hailegiorgis (2014) apply ABM to model cholera outbreaks in Kenyan refugee camps, modeling individuals as distinct agents making decisions regarding their own behavior at each time step. They concluded that cholera spreads radially from contaminated water sources and that seasonal rains may result in cholera outbreaks, thus stressing the importance of environmental reservoirs.

The environmental reservoirs spread cholera bacteria between different regions. Bertuzzo et al. (2008) investigated the spreading of cholera epidemics by explicitly modeling the river networks. Each region is represented using a local compartmental model described by the SIB model presented in Codeço (2001). However, the bacteria concentration in a region would include a transmission rate between regions. The dispersal rate of bacteria between regions depends on the flow direction of the river network, the degree of each node in the river network and the aquatic reservoir size for each region. Mari et al.

(2012) employ a similar modeling scheme, but include long-distance bacteria dispersal through human mobility, by introducing a probability of migration based on population in each region and the distance between the regions. The spatially explicit cholera model is applied to the 2010 Haiti cholera epidemic in Bertuzzo et al. (2011). Here, human mobility is included, but instead of adding an additional term to the bacteria concentration department, the migration dispersal is implicitly modeled by redefining the probability of dispersal. Instead of focusing solely on river networks, Bertuzzo et al. (2011) account for the distance between regions and their populations.

Environmental variability can be introduced to better account for the uncertainty. Eisenberg et al. (2013) examined the relationship between rainfall and cholera outbreaks in Haiti, and concluded that increased rainfall is significantly correlated with increased cholera incidence for up to a week later. The importance of accounting for environmental variability is also emphasized in (Allen, 2017), particularly for waterborne diseases, such as cholera. King et al. (2008) incorporated the environmental fluctuations in cholera models by including a Gaussian white noise term to the infection rate. Azaele et al. (2010) employed a similar approach, including a delta-autocorrelated Gaussian noise term for the infected compartment, while Gazi et al. (2010) accounted for environmental fluctuations by including Gaussian white noise for every compartment. Allen (2016) compared the white noise approach to mean-reverting processes for a generic epidemic model, and argued that mean-reverting processes are more biologically plausible because they cannot drift towards infinity, but will instead move back to their asymptotic mean.

Control and intervention strategies can be incorporated in epidemic models. Certain interventions can be incorporated as distinct compartments. Liu et al. (2019) proposed a Susceptible-Infected-Quarantined-Recovered (SIQR) cholera model, where infected individuals would get quarantined with a certain rate. Although quarantine can be effective against certain contagious diseases, the authors note that it is controversial due to interfering with individual rights, and WHO has emphasized that it is unnecessary if it may divert resources from other interventions (World Health Organization, 2010). Mwasa and Tchuenche (2011) developed a cholera model that included compartments for educated individuals, vaccinated individuals and treated individuals, in addition to quarantined individuals. Instead of defining a distinct vaccinated compartment, vaccinations may also be modeled using the recovered compartment, as done in Tuite et al. (2011) and Wang and Modnak (2011). However, the immunity duration of recovered and vaccinated individuals may differ, thus requiring distinct compartments. An example of such an approach is Andrews and Basu (2011), which assumes that a recovered individual will, on average, lose its immunity after 10 months, while a vaccinated individual will, on average, be immune for 2 years.

Interventions can also be modeled as dynamic parameters. Wang and Modnak (2011) included vaccination, therapeutic treatment and water sanitation. Vaccinated individuals

transitioned to the recovered compartment. Therapeutic treatment is modeled as a parameter that increases the rate of recovery. Water sanitation is modeled as a death rate for the bacteria concentration department. The vaccine efficacy and the exact treatment given is not specified. Andrews and Basu (2011) model similar interventions: vaccination, clean water and antibiotic treatment. They accounted for vaccine efficacy, that is, not all vaccinated individuals developed immunity. Instead of modeling water sanitation as increasing the death rate of cholera bacteria, Andrews and Basu (2011) modeled the intervention as a reduction in the susceptible population consuming contaminated water. Both approaches are realistic and reasonable. An increase in death rate is reasonable when the water is centrally disinfected, while a decrease in consumption of contaminated water might be more appropriate when households are given disinfectant to clean their own water. Lastly, Andrews and Basu (2011) modeled a specific treatment, antibiotics, which they argued would both increase the recovery rate and decrease the excrete rate of symptomatic individuals.

3.2 Emergency Logistics Overview

Emergency logistics spans a wide variety of problems. Different types of emergencies requires varying responses, and are thus distinct problems. What all these problems have in common, in contrast to the more traditional and heavily researched business logistics, is the chaotic situation, possibly compromising the flow of information. While demand is uncertain in a business logistics setting, the surge of demand following an emergency and the disarray of the situation, entitles the need of research specifically on emergency logistics.

Altay and Green III (2006) first surveyed operation research (OR) and management sciences (MS) research on disaster operations management (DOM). Here, they identified that during emergency situations, the duration and scale is uncertain, the problem environment is chaotic and may change rapidly, and decisions must be taken promptly with little or no, and possibly unreliable, data. Having identified how decision-making during disasters differ from conventional decision-making, they set the stage for a wide variety OR and MS research within the emergency and disaster logistics and management field, some of which can be applied in epidemic logistics.

Emergencies can be divided into two phases: pre-disaster and post-disaster, which in turn are typically divided into four activities: mitigation and preparedness, and response and recovery (Coppola, 2006). The pre-disaster situation requires planning and consists of mitigation and preparation of the emergency, involving decisions such as the location of medical distribution centers and stock pre-positioning of medical supplies. The post-disaster consists of response and recovery after the emergency has occurred. Relevant decisions include effective and fair distribution of medical supplies and transportation of

victims and casualties.

Caunhye et al. (2012) classified emergency logistics into three problem categories: facility location, relief distribution and casualty transportation. The facility location focuses on the location of various types of facilities, typically to maximize coverage or minimize response time. During mitigation, preparation, response and recovery, the location and inventory of different types of facilities should be determined. The different facilities to locate in all four emergency activities are shown in Figure 3.2, inspired by a similar figure in Boonmee et al. (2017).

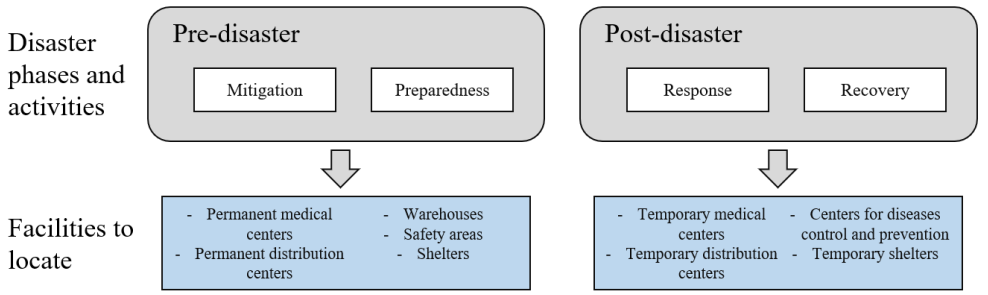

Figure 3.2: Facility location decisions at different stages of an emergency, modified from Boonmee et al. (2017).

Relief distribution concerns bringing relief, for instance, medical supplies, to the individuals affected by the emergency. The problem of relief distribution can be approached as a commodity flow problem, deciding what relief to be distributed across which routes. Another approach to relief distribution is resource allocation, where the flow quantity across routes is not determined. Instead, the focus is to allocate supplies at facilities to best provide relief to the affected individuals and areas. While relief distribution concerns bringing relief out to the affected areas, casualty transportation regards the transportation of affected individuals from the emergency areas to safety, for instance, medical treatment centers.

Several works combine the pre-disaster and post-disaster decisions. Mete and Zabinsky (2010) developed a two-stage stochastic programming model combining facility location and relief distribution. The first stage selected warehouse locations and inventory levels at the warehouses. After the disaster scenario is observed and the post-disaster response phase is entered, the recourse decisions are the amount of medical supplies to be transported from a warehouse to a hospital. Rawls and Turnquist (2010) also developed a two-stage stochastic programming model, combining preparedness and response and focusing on hurricane emergencies. Similar to Mete and Zabinsky (2010), they have facility location and pre-positioning of materials as first-stage decisions. In addition, Rawls and

Turnquist (2010) included material destruction as a stochastic element in different scenarios. The recourse decisions are thus to distribute the remaining pre-positioned supplies to meet demand. The trajectory of a hurricane is uncertain, making it possible for several disasters at different locations over time. Rawls and Turnquist (2012) extended the two-stage model to a multi-stage model with chance constraints for demand coverage, thus accounting for the several possible trajectories of the hurricane.

3.3 Epidemic logistics

Different emergencies have different characteristics, and as emphasized in Gupta et al. (2016), the models should account for these differences and be cautious claiming they work for any emergency or disaster type. After an earthquake has occurred, no or minor earthquake events arise later on, it does not evolve like the spread of a disease during an epidemic or the trajectory of a hurricane. Epidemics can also be contained using intervention methods, while hurricanes are impossible to stop, thus requiring relief in the recovery phase, not containment interventions in the response phase. This section focuses on emergency logistics for epidemic outbreaks. First, the various problem formulations and methodologies are identified in Section 3.3.1. Then, different objectives are investigated in Section 3.3.2. In Section 3.3.3, the various interventions and resources allocated in the existing literature are identified. Lastly, Section 3.3.4 identifies previous works combining epidemic models with epidemic logistics models.

3.3.1 Methodologies

Becker and Starczak (1997) developed a deterministic linear program for optimal vaccine allocation for households in a community. Tanner et al. (2008) extends the work by introducing chance constraints, that is, ensuring the constraints are satisfied with a given probability. Tanner and Ntaimo (2010) developed a branch-and-cut algorithm and applied it on the chance-constrained optimal vaccine allocation problem proposed in Tanner et al. (2008). Although the resources allocated during outbreaks are discrete, several works simplify the integer constraint. Arora et al. (2010) allocates ratios of the available resources and Yarmand et al. (2014) relax the integer constraints, allowing for continuous number of vaccines to be allocated. Most models include integer variables for the resources, resulting in mixed-integer linear programs (MILP), such as Liu et al. (2015), Anparasan and Lejeune (2017) and Büyüktahtakın et al. (2018), with integer medical resources, medical personnel and treatment centers, respectively.

Certain works also include nonlinear formulations. Büyüktahtakın et al. (2018) proposed a mixed-integer nonlinear program (MINLP). However, the nonlinearity arose from a minimum-constraint and the constraint was linearized, resulting in a MILP formulation.

Wang et al. (2009) proposed a MINLP for optimal material distribution with a nonlinear objective. Arora et al. (2010) presented a quadratic program for resource allocation during pandemics.

Anparasan and Lejeune (2017) and Arora et al. (2010) both propose static, single-period models for resource allocation during epidemic outbreaks. However, as emphasized in Arora et al. (2010); epidemic spread is inherently dynamic. Therefore, a wide range of literature focus on dynamic formulations with multiple allocation periods. Zaric and Brandeau (2001) and Zaric and Brandeau (2002) both proposed models allocating intervention investments over short time horizons. Rottkemper et al. (2012) proposed a rolling-horizon model for inventory relocation. Yaesoubi and Cohen (2011), Coşgun and Büyüktahtakın (2018) and Long et al. (2018) all proposed approximate dynamic programming (ADP) methods to solve resource allocation problems for dynamic epidemic outbreaks.

3.3.2 Model objectives

In contrast to business logistics, the objective in emergency logistics and humanitarian operations is not necessarily to minimize cost or maximize profits. Kovacs and Moshtari (2019) identified the need for objectives aligning with real stakeholders in humanitarian operations literature, focusing on saving lives and minimizing damages within a budget constraint, not minimizing costs.

There are three broad variations in model objectives. First, several works apply the traditional business logistics objective with minimization of costs. Examples include Wang et al. (2009) and Liu et al. (2015), both minimizing the transportation costs. Others incorporate fatalities and human suffering in a cost-based objective, thus employing a net monetary benefit approach. Zaric and Brandeau (2001), Zaric and Brandeau (2002) and Brandeau et al. (2003) all used a quality-adjusted life years (QALY) objective, where each epidemic compartment is assigned a quality, and the quality is maximized, summing over all time periods and compartments, but employing a discount rate for later periods. Yaesoubi and Cohen (2011) also had a net monetary benefit objective, maximizing the difference between the decision-makers willingness to pay for health interventions and the expected costs of implementing the intervention given the future disease spread. Both Ludkovski and Niemi (2010) and Yarmand et al. (2014) minimized intervention costs, but Ludkovski and Niemi (2010) also included the cost of having infected individuals, thus accounting for social cost.

Several recent works have focused on objectives other than variations of cost-minimization, such as maximizing demand coverage, that is, assist as many casualties as possible, or minimizing infections or fatalities. Rachaniotis et al. (2012), Büyüktahtakın et al. (2018) and Long et al. (2018) minimized total number of new infections, while Ren

et al. (2013) minimized total number of fatalities. Coşgun and Büyüktahtakın (2018) presented a multi-objective formulation, minimizing the weighted sum of HIV-infections, AIDS-infections and fatalities. Anparasan and Lejeune (2017) also aligned their model objective with real stakeholders, but instead of minimizing infections or fatalities, they maximized demand coverage, that is, how many cholera infected individuals they treated.

Several works have also investigated the relationships among various objectives. Brandeau et al. (2003) compared the QALY objective with minimizing number of infections, and proved that allocations made greedy in cost-effectiveness ratios and allocations made greedy in infection growth provide different results. Rottkemper et al. (2012) studied the trade-off between unsatisfied demand and operational costs, and found that in some cases the unsatisfied relief demand could be reduced drastically, while only increasing the operational costs slightly.

3.3.3 Interventions and resources

Different approaches to model interventions and resources include generic resources, such as investments and funds (Zaric and Brandeau, 2001, 2002), and an unspecified medical resource (Liu et al., 2015). Others also include a single intervention method, but model a more specific resource. For instance, Arora et al. (2010) allocates antivirals for pandemic relief and Anparasan and Lejeune (2017) allocates medical personnel for ambulance transportation of cholera-infected individuals. Hospital beds is used in several works, such as Büyüktahtakın et al. (2018) and Long et al. (2018). Coşgun and Büyüktahtakın (2018) incorporates several joint interventions using a single budget allocation variable and a binary variable for each mix of interventions employed. Some models include single interventions, but allow for different strategies, and thus different effects. Ren et al. (2013) allocate vaccines, but differentiates between vaccines allocated for a ring vaccination strategy and vaccines allocated for a mass vaccination strategy.

Yaesoubi and Cohen (2011) differentiates between medical and transmission reducing interventions, where the medical interventions includes both vaccines and antiviral drugs, while the transmission reducing interventions includes measures preventive measures such as social distancing and mask use. Ludkovski and Niemi (2010) also includes multiple intervention methods, focusing on isolation and vaccines.

3.3.4 Combining epidemic and operations research models

Although Arora et al. (2010) assumed a static disease environment, they stressed that the decisions made could affect future epidemic spread, and thus the uncertainties and dynamics of disease diffusion should be accounted for. Dasaklis et al. (2012) surveyed epidemics control and logistics literature, and called for a more multidisciplinary field, drawing expertise not only form logisticians, but from epidemiologists as well.

Zaric and Brandeau (2001), Zaric and Brandeau (2002) and Brandeau et al. (2003) were early in including epidemiological modeling in mathematical programming models. They focus on a generic disease, using a SIR epidemic model and a set of general interventions, each with an associated effect on the epidemic parameters. Brandeau et al. (2003) combines epidemic modeling with mathematical programming for multiple regions, although the regions are independent, i.e. does not interact with each other.

Later works have applied the models on a specific disease, but kept the SIR model (Ludkovski and Niemi, 2010; Yaesoubi and Cohen, 2011; Rachaniotis et al., 2012; Ren et al., 2013; Long et al., 2018). Wang et al. (2009) proposed a model for a general epidemic setting, but included an exposed compartment representing the latent period, i.e. they employed a SEIR-model. In addition, they modeled the SEIR-model as a delay differential equation (DDE). Yarmand et al. (2014) and Liu et al. (2015) also included an exposed compartment in their epidemic modeling.

There is limited existing literature on including disease-specific mechanics in the epidemic modeling combined with an mathematical program. Coşgun and Büyüktahtakın (2018) developed a Susceptible-Infected-AIDS-Death (SIAD) model for HIV response and Büyüktahtakın et al. (2018) employed a Susceptible-Infected-Treated-Recovered-Funeral-Buried (SITR-FB) model for Ebola outbreaks. Long et al. (2018) also investigated Ebola outbreaks, but did not include the treated, funeral and buried compartments. Instead, they proposed a SIR-model that included transmission across regions, modeling human mobility as a possible dispersal method.

Yaesoubi and Cohen (2011) propose a dynamic health policy model for influenza using medical treatment, such as vaccination and antiviral drugs, and transmission reduction techniques, such as face masks and social distancing. They apply an ADP model using a stationary discrete-time Markov chain (DTMC). Instead of assuming that the epidemic states are known in advance, i.e. are stationary, Coşgun and Büyüktahtakın (2018) develops an integrated stochastic compartmental model and ADP model, using a Markov chain model with non-stationary transition probabilities. The model allocates a limited intervention budget among HIV disease compartments to minimize the amount of HIV-infected.

The epidemic models can introduce challenges arising from non-linearities in objectives and constraints and no closed-form evaluations of objectives, and several works rely on approximations to tackle these problems. Büyüktahtakın et al. (2018) propose a deterministic MILP for resource allocation, but investigates the problem over a multi-period horizon. Most epidemic-logistics model combining epidemic modeling with mathematical programming propose two separate models, where the epidemic model output are parameters in the mathematical programming model. Büyüktahtakın et al. (2018) combine the two into a single model at the expense of simplifying the epidemic model. To avoid

non-linearities, they assume the transition rate from susceptible to infected is independent of the size of the susceptible population. Ren et al. (2013) also approximate an epidemic model using Taylor expansions and assuming a constant transmission rate to develop a closed-form solution for the number of infections, to include it in their mixed-integer programming problem (MIP). In their myopic model, Long et al. (2018) approximated their epidemic model by turning the differential equations into difference equations with one-week time intervals. Linearizing the epidemic model is sufficient when having one intervention method, but by introducing several intervention methods, each affecting the different compartments in the epidemic model, the implicit effect of an intervention is not captured if linearized, thus possibly resulting in ignoring certain interventions.

Coşgun and Büyüktahtakın (2018) include an epidemic model into their optimization formulation, but does not make assumptions to ensure linearity. Instead, they apply a dynamic programming approach. They allocate budgetary resources across compartments instead of regions, i.e. they investigate preventive, pre-emptive or reactive strategies, and concluded that preventive measures are favorable to antiviral treatment.

3.4 Literature Review Summary

In light of the COVID-19 pandemic the past months, the field of epidemiology has gained renewed attention, giving rise to novel examples of combining epidemiology with other fields, such as machine learning. For instance, Dandekar and Barbastathis (2020) applied a SEIR epidemic model with a neural network to extrapolate the effect of quarantine measures from public data. Although the growing literature makes it more difficult to review all current knowledge, the review performed throughout this chapter should be sufficient to identify gaps in the literature and where future research can provide the most impact.

Section 3.1 reviewed the literature on epidemic modeling. Providing a brief historic perspective on modeling epidemics using differential equations, the focus was on cholera models. Codeço (2001) proposed an SIB-model, thus incorporating the environmental transmission factor in the SIR-model. Her works have later been extended. Bertuzzo et al. (2011) proposed a spatial model for different regions, taking bacteria dispersal through both river networks and human mobility into account. Several approaches to intervention modeling were reviewed. Andrews and Basu (2011) modeled vaccination, clean water and antibiotic treatment, being three important intervention methods for cholera epidemics.

In Section 3.2, an overview of the emergency logistics were provided. Important considerations to make when developing emergency logistics is to identify if the model is pre-disaster, post-disaster or both, if the focus is on emergency mitigation, preparedness,

response or recovery, and what characteristics the specific emergency to be modeled have. Different emergencies have differing properties and thus requires distinct models.

Lastly, Section 3.3 reviewed emergency logistics specifically for epidemic outbreaks. Different methodologies and problem formulations were reviewed. Several works with various objectives was identified. An important difference between business logistics and emergency logistics is that the former will typically minimize cost or maximize profits, while the latter focus on saving lives, minimize casualties or maximize demand coverage. Various intervention methods were also reviewed. Most literature, including recent works, typically include a single intervention, either specific or general, when responding to epidemics. Rachaniotis et al. (2012) also employed a single-resource approach, but mentioned multi-resource as an important direction of future research.

There are several epidemic control and logistics approaches that address specific challenges within the field, including spatial epidemic models, non-stationary transition rates and operational constraints. For instance, both Büyüktahtakın et al. (2018) and Long et al. (2018) address the need for explicitly considering the spatial spread of an epidemic logistics model, that is, geographically varying transition rates. Coşgun and Büyüktahtakın (2018) include non-stationary transition probabilities in their compartmental model, but consider only a single population. Anparasan and Lejeune (2017) address the issues of realistic operational constraints, with a thorough background and a realistic cholera case-study, but consider a static situation. Long et al. (2018) employ an ADP strategy to minimize the total number of infections during an Ebola outbreak, but allocate only one resource and does not take the stochasticity of transmission across regions into account.

Table 3.1 provides a summary of recent, relevant literature for epidemic response and resource allocation. Explanations for the abbreviations used are included in the nomenclature in Table 3.2. This thesis' contributions to the OR and epidemiology literature are:

(i) **A stochastic dynamic programming model of the epidemic outbreak response and control problem.** The model is developed for multi-stage problems to meet the actual decision-making cycles during epidemics. The model can be used to compare policies in advance of outbreaks or in real-time by decision-makers responding to outbreaks, by supplying the available medical resources, demographic data of the region and the initial or current number of infected, vaccinated and recovered individuals, and assumed bacteria concentration in the aquatic reservoir. The model minimizes the total number of cholera-induced fatalities throughout the finite planning horizon of the outbreak by allocating vaccines, rehydration solutions, antibiotics and disinfectants to regions with current outbreaks and regions with high risk of outbreaks.

(ii) **A cholera-specific epidemic model, combining and extending the works of**

Bertuzzo et al. (2011) and Andrews and Basu (2011). The epidemic model accounts for dispersal of bacteria across regions from both human mobility and river networks, and models four intervention methods: vaccines, disinfectants, rehydration solutions and antibiotics. To the best of my knowledge, it is the first time a cholera-specific model has been combined with a mathematical programming model, and the first time all four interventions are included in an epidemic model. It is also, to the best of my knowledge, the first proposed multi-resource and multi-region resource allocation model combined with an epidemic model that is solved using an approximate dynamic programming approach.

(iii) **The resource allocation model is used to find optimal intervention strategies for outbreaks similar to the 2010 cholera outbreak in Haiti.** Note that the model can be extended to similar situations, that is, epidemic outbreaks with limited intervention resources. The compartmental model in this thesis is based on cholera, but can be replaced with other compartmental models for other diseases.

Table 3.1: Selected literature on epidemic response and resource allocation.

Article	Disease	Objective	Methodology	Epidemic model	Stochastic	Multiple intervention resources	Multiple outbreak regions
This thesis	Cholera	Fatalities	ADP	SAIR-B	Yes	Yes	Yes
Büyüktahtakın et al. (2018)	Ebola	Infections	MILP	SITR-FB	No	No	Yes
Coşgun and Büyüktahtakın (2018)	HIV	Infections and fatalities	ADP	SIAD	Yes	No	No
Long et al. (2018)	Ebola	Infections	ADP	SIR	No	No	Yes
Anparasan and Lejeune (2017)	Cholera	Demand coverage	ILP	NA	No	No	Yes
Liu et al. (2015)	Influenza	Transportation cost	MILP	SEIR	No	No	No
Yarmand et al. (2014)	Influenza	Operational cost	SLP	SEIR	Yes	No	Yes
Ren et al. (2013)	Smallpox	Fatalities	MINLP	SIR	No	No	Yes
Rachaniotis et al. (2012)	Influenza	Infections	MILP	SIR	No	No	Yes
Yaesoubi and Cohen (2011)	Influenza	Net monetary benefit	ADP	SIR	Yes	Yes	No
Ludkovski and Niemi (2010)	Influenza	Operational and social cost	RMC	SIR	Yes	Yes	No
Wang et al. (2009)	General	Transportation cost and time	MINLP	SEIR	Yes	No	No

Table 3.2: Nomenclature for literature review summary.

Abbreviation	Name
ADP	Approximate Dynamic Programming
ILP	Integer Linear Programming
MILP	Mixed-Integer Linear Programming
MINLP	Mixed-Integer Nonlinear Programming
NA	Not Applicable
RMC	Regression Monte Carlo
SAIR-B	Susceptible-Asymptomatic-Infected-Recovered-Bacteria
SIR	Susceptible-Infected-Recovered
SEIR	Susceptible-Exposed-Infected-Recovered
SIAD	Susceptible-Infected-AIDS-Death
SITR-FB	Susceptible-Infected-Treated-Recovered-Funeral-Buried
SLP	Stochastic Linear Programming

Chapter 4

Theory

This chapter presents the underlying theory required to understand the models developed in this thesis. Section 4.1 describes methods for predicting evolution of an epidemic. In Section 4.2, Markov decision processes (MDP), a mathematical framework for representing the state, possible actions and the corresponding transitions, is presented. Lastly, Section 4.3 describes approximate dynamic programming (ADP) and techniques to efficiently compute solutions leveraging the MDP framework.

4.1 Epidemic Modeling

Compartmental models in epidemiology consists of different compartments, i.e. homogeneous populations. Each individual in a compartment have the same transition probability, that is, the same probability of transitioning from their current compartment to another.

The following model is a simple, deterministic compartmental model, based on the explanation in Allen (2017). Let $S(t)$, $I(t)$ and $R(t)$ denote the susceptible, infected and recovered population, respectively. Further, assume no births or deaths and let the total population size be $N(t) = S(t) + I(t) + R(t)$. Then a disease with infection rate β and

recovery rate γ can be described by the following system of differential equations:

$$\frac{dS}{dt} = -\beta I \frac{S}{N} \tag{4.1}$$

$$\frac{dI}{dt} = \beta I \frac{S}{N} - \gamma I \tag{4.2}$$

$$\frac{dR}{dt} = \gamma I \tag{4.3}$$

Given an initial condition $S(0)$, $I(0)$ and $R(0)$, the system of differential equations in Equations 4.1-4.3 can be solved numerically.

The model defined by Equations 4.1-4.3 could also be described in terms of transition probabilities, instead of rates. That is, in each time period, each individual can transition with a certain probability. However, when the population size is large, the number of transitions per time period are close to its expected value, due to the law of large numbers, resulting in an epidemic evolution similar to the deterministic model, described by the system of differential equations.

4.2 Markov Decision Processes

The following sections presents a mathematical representation framework for sequential decision problems where the decision is only dependent on the current state, not its previously visited states.

4.2.1 Markov property

A stochastic process X_t has the Markov property if a consecutive state is only conditionally dependent on the current state. Formally, as described in Pinsky and Karlin (2010),

$$\mathbb{P}(X_{t+1} = j \mid X_0 = i_0, X_1 = i_1, \ldots, X_t = i_t) = \mathbb{P}(X_{t+1} = j \mid X_t = i_t) \tag{4.4}$$

Although previous decisions impacts the state you end up in, if the process have the Markov property, there is no need for further information than the current state to make future transitions.

4.2.2 System representation

This section describes the basic notation used for mathematically representing a system of states, actions and transitions. Consider a system with discrete and finite state space $\mathcal{S} = (1, 2, \ldots, |\mathcal{S}|)$. The system changes when an action, also called decision, is taken.

At time t, let the feasible action space be denoted $\mathcal{X}_t$, and let a feasible action be denoted x_t, where $x_t \in \mathcal{X}_t$.

Given a specific state S_t and an action x_t at time t, the probability of transitioning to state $S_{t+1} = s'$ is denoted $\mathbb{P}(S_{t+1} = s' \mid S_t, x_t)$, and henceforth referred to as the transition probability. Observe that the transition probability follows the Markov property, because the transition probability entering state S_{t+1}, depends only on state S_t, not previous states, such as S_{t-1}. When an action x_t is taken in the current state S_t, a cost $C_t(S_t, x_t)$ is incurred. Depending on the problem, a positive cost can be desirable, and is in that case often referred to as reward. The value of the current state S_t is calculated using the value function, denoted V_t, and is evaluated as:

$$V_t(S_t) = \max_{x_t \in \mathcal{X}_t} \big(C_t(S_t, x_t) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(S_{t+1} = s' \mid S_t, x_t) V_{t+1}(s')\big), \tag{4.5}$$

where γ is the discount rate. Equation 4.5 is often referred to as the standard form of Bellman's equation (Powell, 2007, pp. 49).

Because the transitions are stochastic, there is an implicit underlying stochastic variable in Equation 4.5. Denote the stochastic variable as W_t, its realization as ω and its outcome space as Ω_t. It is now possible to define a deterministic transition function:

$$S^M : \mathcal{S}_t \times \mathcal{X}_t \times \Omega_t \to \mathcal{S}_{t+1}, \tag{4.6}$$

where given a state, action and realization, a specific state is $S_{t+1} = S^M(S_t, x_t, W_{t+1})$ (Range, 2019).

A *policy* π is an unambiguous rule of selecting an action x_t given the state S_t. The goal is to develop a policy that maximizes value at the current state, which considers either a finite or infinite amount of future values as well. The optimal policy can be found through backward recursion of Equation 4.5, given an initial condition.

4.3 Approximate Dynamic Programming

Building upon the MDP framework previously presented, this section describes the resulting problems of the framework and possible solutions to efficiently solve an approximation of Bellman's equation.

4.3.1 Curse of dimensionality

The MDP framework can be used to find optimal policies for sequential decision problems. However, the problems can quickly become computationally intractable. Assume that, at time t, the state space is I-dimensional, and that each dimension has L possible values. Similarly, the action space is J-dimensional and each dimension has M outcomes. Lastly, assume the stochastic variable is K-dimensional and can take N different realizations per dimension. While the state space, action space and outcome space are of significant size on their own, the state space at time $t+1$ is immensely huge. From Equation 4.6, the cardinality of this state space can be computed to be: $|\mathcal{S}_{t+1}| = L^I \cdot M^J \cdot N^K$. Even for moderate values of I, J and K, the new state space explodes, making it computationally infeasible to evaluate all possible new states.

4.3.2 Approximate value function

The curse of dimensionality causes the value function described in Equation 4.5 to be computationally intractable even for moderate-sized problems. Instead of solving it with backward recursion, the value function is approximated and then iteratively updated through forward recursion. This way, there is always an estimate for the second term in Equation 4.5, meaning there is no need to enumerate every possible consequence of the decision taken.

The approximation value function is denoted $\bar{V}_t^n$, where t, as before, refer to the time and n refer to the number of iterations for which the value function has been updated. A possible approach for updating the approximate value function would be:

$$\bar{V}_t^n(S_t) = \max_{x_t \in \mathcal{X}_t} \left(C_t(S_t, x_t) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(S_{t+1} = s' \mid S_t, x_t) \bar{V}_{t+1}^{n-1}(s') \right) \tag{4.7}$$

4.3.3 Exogenous information

In some problems, the transition from the current state S_t to another state S_{t+1}, is not deterministic. The transition may also depend on exogenous information that is not available at the time the decision is made. As in Equation 4.6, denote the exogenous information as W_t being realized and first observable for the decision at time t. Further, a possible outcome is denoted a *scenario*, the outcome space consists of all possible outcomes at a time t and is denoted Ω_t and the *sample path* or *realization path* is a path of realizations across several time periods.

4.3.4 Value function approximation representation

The value function approximation $\bar{V}_t^n$ can be represented in different ways. The goal is to provide a reasonable approximation of $V_t^n(S_t)$ (Powell, 2007, pp. 225).

Look-up table: A conceptually simple way to represent $\bar{V}_t^n$ is as a look-up table. That is, for each possible state, there is a mapping to an approximated value. The look-up table representation can quickly become infeasible due to the first curse of dimensionality, a large state space. In addition, if the state space is large, most states may never be visited, even if the number of iterations n is high. Then the performance of the ADP algorithm using $\bar{V}_t^n$ will depend drastically on the initial values of the look-up table. A solution to this is *aggregation*, which updates similar states to the ones visited.

Linear regression: Due to the computationally complexity of representing the value function approximation for high-dimensional state spaces, $\bar{V}_t^n$ can be represented using regression models. The representation of the value function approximation is then defined by:

$$\bar{V}^n(S \mid \beta) = \beta_0^n + \sum_{i \in \mathcal{I}} \beta_i^n f(S), \tag{4.8}$$

where n is not the power, but an index for the nth iteration, that is the regression parameter after n updates. The function $f(\cdot)$ refers to the basis function, extracting the features from the state variable, and $\mathcal{I}$ refers to the features (Powell, 2007, pp. 237). For a more rigorous explanation of linear regression, please see Section 4.4.1.

The linear regression representation of the value function approximation is linear in the parameters, however, not necessarily in the state attributes, because of the basis function. For example, if the value function approximation is quadratic in the available resources R, it can be defined as:

$$\bar{V}(R \mid \beta) = \beta_0 \sum_{a \in \mathcal{A}} (\beta_{1a} R_a + \beta_{2a} R_a^2), \tag{4.9}$$

where R_a refers to the available resources with attribute a. Linear regression can thus handle more than just functions linear in the state variable. However, the precise form of the value function approximation might not be known in advance. When this is the case, an alternative representation of the is using a *neural network*, a mathematical tool explained in the following section.

4.4 Neural Networks

Neural networks (NN) are computing systems composed of interconnected neurons. Each neuron is associated with an activation function and a set of weights. NNs are inspired by biological neural networks, however, they are not designed to be perfect models of the brain. Instead, NNs are used to approximate some mathematical function $f(\cdot)$. The following sections are largely based on a similar description of neural networks in Olstad and Verås (2019).

4.4.1 Linear Regression

A linear regression model aims to model the relationship between one or more explanatory variables, and a real-valued prediction variable. Mathematically, a linear regression model with explanatory variables $\boldsymbol{x} \in \mathbb{R}^n$ takes the form

$$f(\boldsymbol{x}) = \beta_0 + \sum_{i=1}^{n} \beta_i \boldsymbol{x}_i. \tag{4.10}$$

The coefficients β_i are found my minimizing some statistical error. Examples of such error statistics are the mean squared error (MSE), minimizing the square deviation from the actual values, and mean absolute error (MAE), minimizing the absolute deviation from the actual values. MSE and MAE can be mathematically expressed as:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \tag{4.11}$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^{n} \|y_i - \hat{y}_i\| \tag{4.12}$$

Here y_i denotes the actual value of observation i, $\hat{y}_i$ is the predicted value and n is the number of observations. MAE and MSE will be equivalent in the case where the difference between y_i and $\hat{y}_i$ is equal at all i. The main difference is that the MSE penalizes large errors more than the MAE.

4.4.2 Perceptron

The perceptron is a binary classifier, and can be seen as a basic building block of a feed-forward neural network (FFNN), also called Multilayer Perceptrons (MLP). A perceptron computes the scalar product of an input vector $\boldsymbol{x} \in \mathbb{R}^n$ and a weight-vector $\boldsymbol{w} \in \mathbb{R}^n$, adds a bias term $b \in \mathbb{R}$, and transforms it using the Heaviside step function seen in Table 4.1. The entire perceptron can be described as

$$f(\boldsymbol{x}; \boldsymbol{w}, b) = \begin{cases} 1, & \text{if } \boldsymbol{x}^T \boldsymbol{w} + b > 0 \\ 0, & \text{otherwise} \end{cases} \tag{4.13}$$

4.4.3 Feedforward Neural Networks

When the same input vector $\boldsymbol{x}$ is used for several perceptrons, with potentially different weights, the perceptrons constitute what is called a *layer*. When the output from all perceptrons in one layer is used to several other perceptrons in another layer, the result is a multilayer perceptron model (MLP), which is a type of NN.

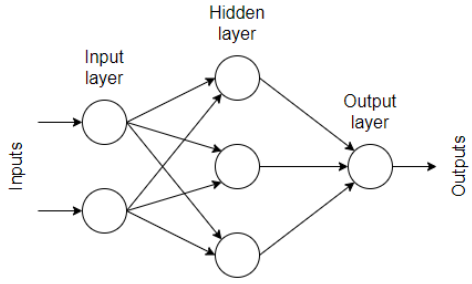

Figure 4.1: Example FFNN with two input neurons, one hidden layer with three neurons and an output layer with one neuron.

When constructing FFNNs, there are several modeling decisions. The selection of appropriate activation, cost and output functions, as well as a functioning architecture and differentiation algorithm is described in the following paragraphs. An example of a FFNN is seen in Figure 4.1.

Activation function: The nonlinear function of the scalar product of the input $\boldsymbol{x}$ and weights $\boldsymbol{w}$ does not need to be the Heaviside step function; it can be any nonlinear function $f\colon \mathbb{R} \to \mathbb{R}$. The perceptron is thus called a *neuron* and the function is called an *activation function*. Typical choices of activation functions include the Heaviside step function, the rectified linear unit function (ReLU), the logistic sigmoid function and the hyperbolic tangent function. These functions are shown in Table 4.1.

Table 4.1: Different activation functions for NNs.

Activation function	Equation
Heaviside step	$f(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases}$
Rectified linear unit	$f(x) = \max(0, x)$
Sigmoid	$\sigma(x) = \frac{1}{1+e^{-x}}$
Hyperbolic tangent	$\tanh(x) = \frac{e^{2x}-1}{e^{2x}+1}$

Cost function: The cost function, also called loss function, defines how to quantify the prediction error for a NN. For regression problems, the cost function are the same as

described for linear regression in Section 4.4.1, that is, MSE and MAE.

Output function: The output function is the activation function of the output layer. For regression problems, a linear output function is common, which is the scalar product of the layer weights and the output from the previous layer is passed on.

Architecture: The architecture of an FFNN refers to the number of hidden layers, the amount of neurons in each layer, and how the layers are connected. The *universal approximation theorem* states that given a wide range of activation functions, including ReLU, a sufficiently large FFNN may approximate any function mapping from any finite discrete space to another (Leshno et al., 1993). However, even if the FFNN may be able to represent the function to be approximated, the NN might not be able to learn it sufficiently fast based on the available data. Thus, the architecture of the NN is an important decision to ensure rapid learning and generalization.

Differentiation algorithm: An FFNN outputs a prediction based on some input features. The input propagates forward through the network, in what is called *forward propagation*. The training of the NN refers to the adjustment of the weights and biases at each layer. A typical approach is to adjust the weights based on the gradient of the cost. To compute the gradient, a differentiation algorithm is needed. The *backpropagation* procedure is extensively used today, and the algorithm recursively applies the chain rule of calculus (Goodfellow et al., 2016).

In a single-layered NN, given an output value and the true target, the cost E is calculated. To update a weight w_{ij} in the network, that is, the weight between the ith neuron of the previous layer and the jth neuron in the current layer, the partial derivative of the error with respect to the specific weight must be computed. Let o_j be the jth output neuron and l_j be the input to the jth output neuron, then Equation 4.14 describes how to recursively apply the chain rule to compute the desired partial derivative:

$$\frac{\partial E}{\partial w_{ij}} = \frac{\partial E}{\partial o_j}\frac{\partial o_j}{\partial w_{ij}} = \frac{\partial E}{\partial o_j}\frac{\partial o_j}{\partial l_j}\frac{\partial l_j}{\partial w_{ij}}. \tag{4.14}$$

Optimization algorithm: The internal optimization problem of an NN is to minimize the generalization error, i.e., the cost function on previously unseen data, by varying the weights and biases. In order to achieve a high performance of NN, large training sets are often necessary. However, large training sets are also computationally expensive. Instead of calculating the gradient based on all records of the training data, the gradient can be calculated by uniformly draw a single data point or a *minibatch*, i.e., a small set of samples from the data. The size of the minibatch typically ranges from a single record to several hundred. Let $L(w)$ denote the cost when using weights w, then the updated

weights are:

$$w^{(k+1)} := w^{(k)} - \gamma \nabla L(w^{(k)}), \tag{4.15}$$

where γ is the *learning rate*, also called step-size, in the direction of the cost gradient $\nabla L(w^{(k)})$, and k is the update index.

To increase efficiency in the optimization, momentum can be included. With momentum, the learning rate increases if the direction of the gradient is the same over consecutive steps. If the gradient points in several different directions, the learning rate is decreased to avoid overshooting the optimum. Adaptive learning is another technique to increase optimization efficiency of the NN, which assigns a learning rate to each parameter in the model, and adjusts the learning rate throughout the training. *Adam* is an optimization algorithm that combines the concept of adaptive learning rates and momentum. It is a widely used optimization algorithm and it is robust to different hyperparameters (Goodfellow et al., 2016, pp. 309).

Chapter 5

Problem Description

In this chapter, the problem of allocating multiple types of medical resources to different regions within a country during the response phase of an epidemic outbreak is described. The problem is applicable to diseases where outbreaks can be anticipated, e.g. due to seasonality, and thus prepared for. It is applicable for outbreaks lasting a few months, and located in regions with limited medical resources, and thus limited availability of supplies, such as vaccines. In addition, the post-disease consequences must be negligible compared to the consequences of being infected, similar to cholera, which can prove fatal within hours unless treated, but typically have no long-term consequences. The modeled disease must also be contagious and it must spread through human migration and environmental mediums, such as rivers. The focus is on allocating the limited resources available at the national level to the country's regions.

When an epidemic outbreak has occurred, the immediate response is establishing temporary medical facilities. The different facilities provide particular sets of interventions. When a facility is established, the region it covers can receive medical resources. As the epidemic progresses, additional medical resources can be allocated.

In epidemic emergencies, when the consequences of being infected are critical and the post-disease consequences are negligible, the overarching objective is to minimize the number of fatalities caused by the outbreaks, given the medical resources available.

To respond to an outbreak, demographic data for the affected and neighboring regions is required and is typically updated periodically with every census the country performs. The information concerning the immediate availability of medical resources and the initial number of infected and vaccinated individuals must be gathered. The demand for medical resources in each region is uncertain and is estimated. Information regarding

the lead times initial stocking and transportation times for the distribution of resources between regions and the resource capacity for various types of facilities are also required.

For resources to be allocated to a region, a facility with available capacity must have been established. The number of medical facilities established in a region is limited by the number of available facility locations for the respective facility types. Establishing a facility requires sufficient available medical personnel. The distribution of resources to regions must account for transportation times, and in addition, certain types of resources can have lead times before being available for distribution at all. Different resources affect future demand in distinct ways, and the resources are limited, and thus possibly not sufficient to properly treat all infected individuals.

Chapter 6

Mathematical Models

In this chapter, the mathematical models used to solve the problem and estimate the demand described in Chapter 5 are presented. Section 6.1 presents a schematic and mathematical description of the compartmental epidemic model. In Section 6.2, the resource allocation model is mathematically formulated and presented.

6.1 Epidemic Model

The epidemic model used for demand estimation can be interpreted as a network of mutually dependent compartmental models for each region, where cholera bacteria are transported through river networks and human migration. The modeling of interventions is inspired by Andrews and Basu (2011) and the modeling of the interaction between regions is based on Bertuzzo et al. (2011). While the epidemic model presented here is designed specifically for cholera, similar models can be developed for other infectious diseases. This can be done by including or excluding compartments, change the transmission mechanics between regions and change the intervention methods, for instance, if there are no available vaccine for the particular disease.

The key assumptions for the epidemic model are:

(i) The rate at which susceptible individuals become infected depends on the contamination of the environment, i.e. the concentration of cholera bacteria in the water reservoir.

(ii) The amount of bacteria dispersal from a region to another decreases with the distance between them and is proportional to the source region's bacteria concentra-

tion times the population size of the destination region (Bertuzzo et al., 2011).

(iii) Recovery from cholera and vaccination provide lasting immunization. While this is not a realistic assumption in endemic areas, it is sufficient for epidemic outbreaks lasting a few months, which are the outbreaks of interest in this thesis. The assumption can be relaxed by introducing immunization fading parameters from the recovered and vaccinated populations to the susceptible population.

(iv) All populations are homogeneous, thus factors such as age and blood type are not taken into account when considering infection and recovery rates. The assumption is widely used in epidemic modeling with compartmental models, but can be relaxed by including additional compartments.

The epidemic model is a set of $6|\mathcal{I}|$ ordinary differential equations, where $|\mathcal{I}|$ is the number of regions included. A region consists of six compartments: susceptible S, asymptomatically infected A, symptomatically infected I, recovered and vaccinated R, cholera-induced fatalities M and *Vibrio cholerae* concentration in water reservoir B. Figure 6.1 shows the epidemic model for a single region, excluding the natural birth and death rates to increase readability.

For each region $i \in \mathcal{I}$, the epidemic model for that region is described by the following system of differential equations:

$$\frac{dS_i}{dt} = \mu(N_i - S_i) - \tau\nu_i - \beta_i \frac{B_i}{\kappa + B_i} S_i \tag{6.1}$$

$$\frac{dA_i}{dt} = p\beta_i \frac{B_i}{\kappa + B_i} S_i - \gamma A_i - \mu A_i \tag{6.2}$$

$$\frac{dI_i}{dt} = (1-p)\beta_i \frac{B_i}{\kappa + B_i} S_i - \mu_C\big(\phi_i + (1-\phi_i)\chi\big) I_i - \gamma\big((1-\theta_i) + \theta_i\lambda\big) I_i - \mu I_i \tag{6.3}$$

$$\frac{dR_i}{dt} = \tau\nu_i + \gamma\Big(A_i + \big((1-\theta_i) + \theta_i\lambda\big) I_i\Big) - \mu R_i \tag{6.4}$$

$$\frac{dM_i}{dt} = \mu_C\big(\phi_i + (1-\phi_i)\chi\big) I_i \tag{6.5}$$

$$\frac{dB_i}{dt} = \frac{\rho_A}{W_i} A + \frac{\rho_I}{W_i}\big(\psi\theta_i + (1-\theta_i)\big) I_i - \mu_B B_i - l\left(B_i - \sum_{j=1}^{|\mathcal{I}|} P_{ij} \frac{W_j}{W_i} B_j\right) \tag{6.6}$$

A definition of all the parameters are found in Table 6.1.

Equation 6.1 describes the changes to the size of the susceptible population in region i. The first term accounts for net birth and death in the region, where $N_i = S_i + A_i + I_i + R_i$, i.e. the total population of the region. The birth and death rates are assumed to be the

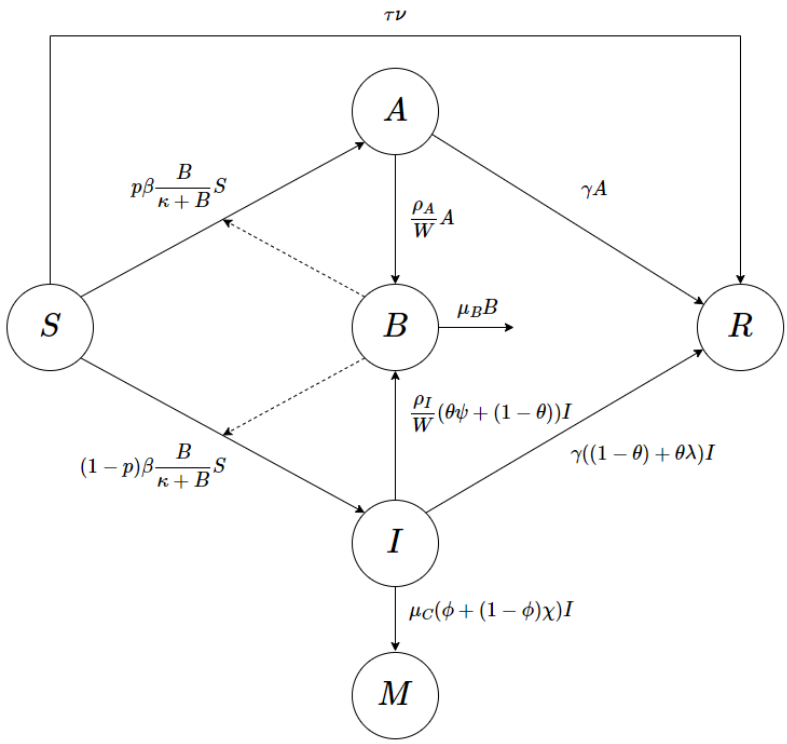

Figure 6.1: Schematic depiction of the epidemic model for a single region with each compartment population and the transition rates among the different compartments. In addition to the rates shown, compartment S has a natural birth rate μN, where N is the population in the region, compartments S, A, I and R have a natural death rate μ and compartment B has transition rates for transmission across regions, which has been discarded in this figure for simplifying purposes.

same and are denoted μ. While this is often not the case, due to the short time horizon of a few months, the differences in the birth and death rates are negligible. The second term in Equation 6.1 captures the vaccination dynamics; τ is the efficacy of the vaccine and ν_i is the absolute number of individuals vaccinated per unit time. The last term refers to the rate of infection, depending on the rate of contaminated water consumption in the region, β_i, and the probability that an individual consuming contaminated water becomes infected. κ is the half-saturation constant, defined as the bacteria ingested for which the probability of becoming infected is 50%.

The changes in the asymptomatic population are described in Equation 6.2. Among the individuals becoming infected, the proportion p of the total population in the region does not develop symptoms, although they can still possibly infect others. The asymptomatic

population recovers with the recovery rate γ and dies of causes other than cholera with the natural death rate μ.

Equation 6.3 describes the changes in the symptomatic infected population. The proportion $(1-p)$ of infected individuals develop symptoms. In region i the proportion ϕ_i of the symptomatic individuals receive rehydration solution. If an individual receives such treatment, the death rate due to cholera is μ_C. However, if no rehydration treatment is received, the cholera-induced death rate is expedited to $\chi\mu_C$. If a symptomatic individual receives antibiotics, the recovery rate increase from the natural rate γ to rate $\lambda\gamma$. The number of recoveries of symptomatic individuals also depends on the current amount of symptomatic individuals and the proportion receiving antibiotics θ_i. As with the other compartments, symptomatic individuals can also die of other causes that cholera with rate μ.

The changes in the recovered and vaccinated population in region i is described by Equation 6.4. The first term describes the transitions from susceptible population getting vaccinated, the second term is recovering individuals from both the symptomatic and asymptomatic populations. The last term describes the number of deaths not related to cholera.

Equation 6.5 describes the increase in cholera-induced fatalities. Transitions are only possible from the symptomatically infected compartment. The equation is not necessary to solve the system of differential equations, but is used in the resource allocation model, and thus included.

The last equation in the system, Equation 6.6, describes the changes in cholera bacteria concentration in the water reservoir in the region. The asymptomatic and symptomatic infected individuals are assumed to excrete bacteria with rate ρ_A and ρ_I, respectively. The first term describes the bacteria excreted to the environment by asymptomatic individuals. The second term describes the same, but for symptomatic individuals, as well as differentiating on whether the symptomatic individuals receive antibiotics or not. Receiving antibiotics are assumed to reduce the bacteria excretion rate with a factor ψ. The concentration of bacteria decays naturally with a net rate μ_B. The last term in Equation 6.6 describes the transmission between regions and is based on Bertuzzo et al. (2011). The bacteria disperse between regions at rate l, and the total dispersal between two regions depends on the current concentration of cholera bacteria in the water sources for the two regions, as well as the relative size of the water sources. Bacteria dispersal from region i to region j occur with the following probability:

$$P_{ij} = \frac{N_j e^{-\frac{d_{ij}}{D}}}{\sum_{k\neq i}^{|\mathcal{I}|} N_k e^{-\frac{d_{ik}}{D}}}, \tag{6.7}$$

where d_{ij} is the dispersal distance between regions i and j, and D is the mean dispersal distance. To decrease the computation time significantly, the population size is assumed constant throughout the epidemic when calculating the bacteria dispersal probability. This is a reasonable assumption, expecting the epidemic is responded to, and that the natural birth and death rates are insignificant due to the relatively short horizon, thereby making the total number of fatalities small compared to the original total population.

Table 6.1: Parameter definitions, values and references.

Parameter	Definition
ν_i	Rate of vaccination
β_i	Proportion of individuals consuming contaminated water
ϕ_i	Proportion of individuals receiving rehydration treatment
θ_i	Proportion of individuals receiving antibiotics
W_i	Water reservoir size
τ	Vaccine efficacy
κ	Half-saturation constant
p	Proportion of infected being asymptomatic
γ	Rate of recovery
ρ_A	Rate of excretion, asymptomatic individuals
ρ_I	Rate of excretion, symptomatic individuals
χ	Relative rate of cholera-induced death, not receiving rehydration treatment
λ	Relative rate of recovery, receiving antibiotics
ψ	Relative rate of excretion, receiving antibiotics
μ	Rate of birth and non-cholera induced death
μ_B	Net rate of cholera bacteria decay
μ_C	Rate of cholera-induced death
l	Mean rate of regional bacteria dispersal

Most parameters in the epidemic model can be set from the literature, such as the half-saturation constant, or be estimated with similar methods from the literature, such as the water reservoir size in a region. Other parameters have been calibrated to historical data, and thus might not be applicable in other regions. The rate of cholera bacteria dispersal between regions l is based on the parameter value in Bertuzzo et al. (2011), where it was calibrated to best fit the cumulative cholera cases. However, the dispersal is likely affected by environmental fluctuations, such as local climatic conditions. To capture the parameter uncertainty and environmental variability, l is modeled as a stochastic param-

eter following some specified distribution. The parameter is assumed to be realized on a weekly basis for each time period. Furthermore, it is assumed given when solving the system of differential equations in Equations 6.1-6.6.

6.2 Resource Allocation Model

The resource allocation model is described using the Markov decision process (MDP) framework. In the following sections, the key assumptions are presented, before the model is formulated to analyze problems of the type described in Chapter 5. The full notation and model is included in Appendix B.

6.2.1 Assumptions

Besides the assumptions given in the problem description, the additional assumptions listed are made to facilitate a mathematical programming model:

(i) The model has a finite planning horizon, assuming the epidemic is eventually eliminated. The planning horizon is divided into a number of discrete time periods, and decisions are made periodically at the beginning of each time period.

(ii) The resource allocation focus is on the national level. The model considers aggregate capacities at the national and regional level and assumes that the decision-makers within the region coordinate the allocation of personnel and intervention resources. A region is also considered a single demand point, thus deployment of intervention resources within a region ignores the intraregion transport times.

(iii) The number of personnel and intervention resources available at the national level is assumed to be given, but intervention resources can arrive throughout the planning horizon. The medical personnel is considered homogeneous. Medical personnel is also assumed to be immune to the disease, and thus the amount of personnel is constant throughout the planning horizon.

(iv) The transportation capacities and times are ignored in the model. Transportation is assumed to be immediately available to the decision-makers or available through local aid form the population.

6.2.2 State variables

Given sets of regions $\mathcal{I}$, medical facility types $\mathcal{N}$ and intervention methods $\mathcal{M}$, the state of the system at time t is described by $S_t = (R_t, D_t, M_t)$, which captures the total supply of medical resources, the demand for different medical resources in the regions and the cumulative cholera-induced fatalities. The state S_t is observed before any decisions are

made in time period t. The resource vector $R_t = (R_{tm})_{m\in\mathcal{M}}$ represents the amount of medical resources allocated for intervention type m available at time t. The vector $D_t = (D_{tim})_{(i,m)\in\mathcal{I}\times\mathcal{M}}$ describes the demand for medical intervention type m in region i at time t. The scalar $M_t = \sum_{i\in\mathcal{I}} M_{ti}$ is the sum of cumulative cholera-induced fatalities over all regions.

6.2.3 Decision variables

The primary decision to be made is the allocation of medical resources to regions. However, the allocation of resources is restricted by the capacity of the established medical facilities and the available personnel in the region. The primary decision variable $x_t = (x_{tim})_{(i,m)\in\mathcal{I}\times\mathcal{M}}$ determines the number of resources for intervention type m allocated and deployed in region i at time t. The supporting decision variable $y_t = (y_{tin})_{(i,n)\in\mathcal{I}\times\mathcal{N}}$ corresponds to the number of facilities of type n open in region i at time t. Lastly, the supporting decision variable $z_t = (z_{ti})_{i\in\mathcal{I}}$ determines the number of personnel allocated to region i at time t.

At the country level, there cannot be deployed more resources than available and resources at the regional level can not exceed the demand:

$$\sum_{i\in\mathcal{I}} x_{tim} \leq R_{tm}, \qquad \forall t \in \mathcal{T}, \forall m \in \mathcal{M} \tag{6.8}$$

$$x_{tim} \leq D_{tim}, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{6.9}$$

In order to deploy resources to region i, there must be sufficient facility capacity to store and apply interventions. Each facility of type n has capacity B_{nm} for intervention type m:

$$x_{tim} \leq \sum_{n\in\mathcal{N}} B_{nm} y_{tin}, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{6.10}$$

Each intervention type m takes the fraction workload U_m during the time period. The total amount of work given the deployment of medical resources must be covered by the available personnel in the region:

$$\sum_{m\in\mathcal{M}} U_m x_{tim} \leq z_{ti}, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I} \tag{6.11}$$

Establishing a medical facility of type n requires a minimum amount of personnel Q_n, thus there must be enough medical personnel allocated to the region to operate all the

established facilities within the time period:

$$\sum_{n \in \mathcal{N}} Q_n y_{tin} \leq z_{ti}, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I} \tag{6.12}$$

The number of medical facilities of type n established in a region is restricted by the availability of locations L_{in}:

$$y_{tin} \leq L_{in}, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I}, \forall n \in \mathcal{N} \tag{6.13}$$

The number of medical personnel allocated to the regions cannot exceed the nationally available amount of medical personnel P:

$$\sum_{i \in \mathcal{I}} z_{ti} \leq P \tag{6.14}$$

Finally, the variables are non-negative integers:

$$x_{tim} \in \mathbb{Z}^+, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{6.15}$$
$$y_{tin} \in \mathbb{Z}^+, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I}, \forall n \in \mathcal{N} \tag{6.16}$$
$$z_{ti} \in \mathbb{Z}^+, \qquad \forall t \in \mathcal{T}, \forall i \in \mathcal{I} \tag{6.17}$$

The set of feasible solutions at time t is denoted $\mathcal{X}_t$, where each solution (x_t, y_t, z_t) satisfies Constraints 6.8-6.17.

6.2.4 Information process

During time period t, the exogenous information ω is revealed. The information impacts the demand and fatalities in the next period. However, the future demand is also dependent on the decisions made in period t. Let W_{t+1} refer to the exogenous information realized during time t and thus first available at time $t+1$, where ω occurs with probability $\mathbb{P}(W_{t+1} = \omega)$. When combined with the epidemic model, the exogenous information is defined to be the bacteria dispersal rate between regions, previously denoted l.

6.2.5 Transition function

The transition function for a state S_t is divided into three distinct transition functions, one for each state component. Transitioning from time t to $t+1$ yields the state $S_{t+1} = (R_{t+1}, D_{t+1}, M_{t+1})$. After making decisions in time period t and after the exogenous

information is realized, the state transitions. The resource transition function, that is, the function determining the effect of applying a decision on the resources, is given by:

$$R_{t+1} = \left(R_{tm} - \sum_{i \in \mathcal{I}} x_{tim} + w_{tm}\right)_{m \in \mathcal{M}}, \tag{6.18}$$

that is, allocated resources are withdrawn and utilized. However, new resources may also arrive. Resources of type m, arriving during time period t are denoted w_{tm} and included in the resources at time $t + 1$.

The demand is highly uncertain and depends both upon the decisions taken within the respective time period and the exogenous information. The demand is estimated using the epidemic model developed in Section 6.1. Let $D^E(\cdot)$ refer to the estimated demand from numerically solving the epidemic model for the next period, then the demand transition function is:

$$D_{t+1} = D^E(D_t, x_t, W_{t+1}) \tag{6.19}$$

The cumulative cholera-induced fatalities transition function increments the current value with the new fatalities occurring over the next time period. The new fatalities within a time period is estimated using the epidemic model developed in Section 6.1. Let $C^E(\cdot)$ be the numerically solved estimated number of lives lost due to cholera within a time period from the epidemic model:

$$M_{t+1} = C^E(S_t, x_t, W_{t+1}) \tag{6.20}$$

6.2.6 Cost function

The cost function refers to the immediate cost or benefit from taking an action x_t. At each time period, the epidemic outbreak cost lives. The number of lives lost due to the outbreak, within a time period, is estimated using the epidemic model. The lives lost depends on, among other things, the bacteria concentration in the water reservoir, and thus include the realization of the exogenous information during time period t and first observed at time $t + 1$, previously defined as W_{t+1}. As defined above, $C^E(\cdot)$ estimates number of lives lost within a time period. Thus the cost function is:

$$C_t = C^E(S_t, x_t, W_{t+1}) \tag{6.21}$$

6.2.7 Objective function

The objective is to minimize the expected number of cholera-induced fatalities throughout the planning horizon. This is achieved by identifying the policy π, i.e. the decision

rule given a state, that minimizes the number of fatalities. Let x_t^π denote the resource allocation decision made at time t, employing policy π. The objective is to find the policy π that minimizes the excepted cholera-induced fatalities over the planning horizon, formally denoted as:

$$\min_{\pi \in \Pi} \mathbb{E}\left\{ \sum_{t \in \mathcal{T}} C_t^E(S_t, x_t^\pi, W_{t+1}) \right\}. \tag{6.22}$$

Chapter 7

Solution Methods

This chapter presents solution methods to solve the problem presented in Chapter 5, using the epidemic model (EM) and resource allocation model (RA) described in Chapter 6. Section 7.1 describes the necessary transformation for the RA and EM to communicate. Section 7.2 describes the approximate dynamic programming (ADP) method and its central components, which is the overarching solution algorithm. In this approach, the resources are allocated for all regions at each time period, and the solution algorithm is referred to as stage decomposition. Section 7.3 presents strategies to solve the subproblem arising from the ADP approach. Lastly, Section 7.4 describes an alternative to the stage decomposition solution scheme: a regional decomposition approach, where the solution algorithm is applied for each region individually.

The structural difference between stage and regional decomposition is illustrated in Figure 7.1. The medical resources need to be distributed over regions and over time. The stage decomposition approach divides the planning horizon into time periods, and allocate medical resources to all regions for each time period. The regional decomposition aims to decrease run time by allowing parallelization by region, but possibly at the cost of performance. The approach divides available resources among the regions, before each region employs a stage decomposition strategy for their allocated resources.

A flowchart of the solution algorithm is illustrated in Figure 7.2. The epidemic model is used to supply the resource allocation model with demand parameters and to calculate its objective. The value function approximation (VFA) is updated iteratively in the ADP framework. After a predetermined number of updates, the problem is solved using the VFA. With stage decomposition, the solution method shown in Figure 7.2 is run at the national level. With regional decomposition, there is one such model for each region.

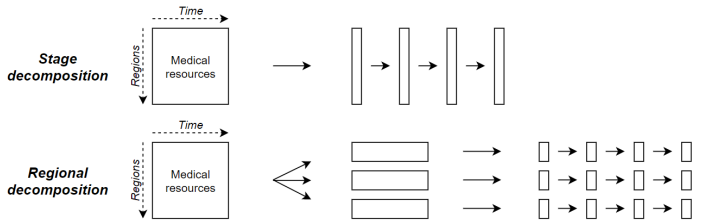

Figure 7.1: Illustration of stage and regional decomposition. Stage decomposition allocates resources to all regions at each time period. Regional decomposition allocates the resources at the beginning, and let each region decide at which time period to employ the resources.

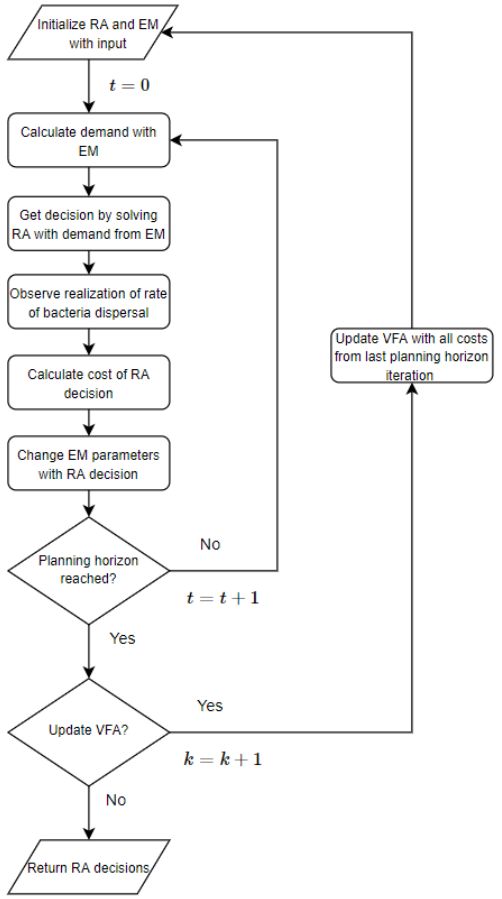

Figure 7.2: Flowchart of solution algorithm using decomposition by stage. Abbreviations used: Resource allocation model (RA), Epidemic model (EM), Value function approximation (VFA).

7.1 Epidemic and Resource Allocation Model Interaction

When the resource allocation model and the epidemic model interacts with each other, the compartmental populations in the epidemic model are transformed into demand for the resource allocation model, and the decisions made by the resource allocation model are translated into parameters for the epidemic model. These transformations are explained in the following paragraphs.

7.1.1 Demand transformation

The compartments in the epidemic model provides demand parameters as input to the resource allocation model. Which compartments constitute demand depends on the intervention method. Vaccines and disinfectants are applied to the susceptible compartment, and therefore the number of susceptible individuals in a region is defined as the demand for vaccines and disinfectants in the respective region. Similarly, rehydration treatment and antibiotics are applied to symptomatic infected individuals, hence the demand in the region corresponds to the number of symptomatic infected in the region. Resources are sent in kits, and therefore the demand is ceiled, i.e. rounded up. For instance, if there are 800 infections in a region and the kit size is 1000, the demand for rehydration solutions and antibiotics is 1.

7.1.2 Decision transformation

The decisions made in the RA model determine the number of resources, the number of medical personnel and the number of medical facilities to establish in each region. The resources sent impact the future evolution of the epidemic model by changing the parameters in the epidemic model. After the decisions are made in the resource allocation model, the decision-dependent parameters in the epidemic model are immediately changed. Although this is not practically reasonable, the resources would, nevertheless, be distributed shortly after the decision is made, thus making the effects of the assumption negligible.

With the exception of vaccination, the decision-dependent rates in the epidemic model are defined in terms of the proportions receiving treatment. For instance ϕ_{it} is the proportion of symptomatic infected individuals receiving rehydration solutions in region i at time t. Because the resources last one week and are sent in kits, the resources sent can exceed the actual number of infections. If this is not adjusted for, the proportion of symptomatic infected individuals receiving rehydration solutions may exceed 1, which is not reasonable practically. The demand can also increase throughout the week, reducing the proportion. The rates are set daily to ensure numerical stability. Let q be the kit size, $I_i(t)$ the number of symptomatic infections in region i at time t, and the number of ORS

kits sent to region i for time t, $x_{ti,ORS}$, the rehydration treatment rate is:

$$\phi_{it} = \begin{cases} \max(q \cdot x_{ti,ORS}, 1), & \text{if } I_i(t) \neq 0 \\ 0, & \text{otherwise} \end{cases} \tag{7.1}$$

7.2 Stage Decomposition

This section presents the ADP framework used to temporally decompose the problem. First, the original objective is reformulated in Section 7.2.1. Then, the value function approximation learning strategy is presented in Section 7.2.2. Lastly, Section 7.2.3 describes the VFA modeling choices.

7.2.1 Objective reformulation

The optimization problem presented in Section 6.2 concerns minimizing the number of cholera-induced fatalities throughout the epidemic outbreak. The optimal solution can be found using dynamic programming, which is more clear when the objective, stated in Equation 6.22, is reformulated to Bellman's equation, with stochastic immediate costs. The resulting equation can be solved recursively:

$$V_t(S_t) = \min_{x_t \in \mathcal{X}_t} \mathbb{E}\big\{C^E(S_t, x_t, W_{t+1}) + V_{t+1}(S_{t+1})|S_t\big\}. \tag{7.2}$$

Finding an exact solution requires enumerating every possible permutation of decisions, states and outcomes over time, which, due to the curse of dimensionality, quickly becomes computationally infeasible. Instead of directly calculating the expected cost of being in state S_{t+1}, the future cost is approximated, that is, the true value function $V_t(\cdot)$ is approximated by $\bar{V}_t(\cdot)$. Thus, at each time period, the cost of transitioning from state S_t with decision x_t and the future expected incurred cost of making x_t is accounted for. The objective is reformulated to:

$$\bar{V}_t^k(S_t) = \min_{x_t \in \mathcal{X}_t} \mathbb{E}\big\{C^E(S_t, x_t, W_{t+1}) + \bar{V}_{t+1}^{k-1}(S_{t+1})|S_t\big\}. \tag{7.3}$$

The optimization is now reduced to sequentially solving a resource allocation problem for each time period, where the consequences for future decisions are penalized with the VFA $\bar{V}_t^k(\cdot)$. The index k refer to the number of times the VFA has been updated. The ADP solution algorithm is based on a value iteration approach and described in Algorithm 1.

Algorithm 1 determines a set of decisions, one decision for each time period t, given the planning horizon to allocate for, the initially available resources and immediate demand,

Algorithm 1: StageDecomposition(S_0, $\mathcal{T}$, K)

Input : Initial state S_0,
Planning horizon $\mathcal{T}$,
Value function update iterations K.

Output: Best found decisions $(x_t)_{t \in \mathcal{T}}$.

for $k = 1, \ldots, K$ **do**
  Initialize state path, $StatePath \leftarrow []$;
  Initialize cost path, $CostPath \leftarrow []$;
  Draw realization path,
  $(W_1, W_2, \ldots, W_{\mathcal{T}+1}) \sim$ ExogenousProbabilityDistribution();
  for $t = 0, \ldots, \mathcal{T}$ **do**
    Get decision, $x_t \leftarrow$ TrainingDecision(S_t);
    Observe exogenous information from realization path, $\omega \leftarrow \mathrm{W}_{t+1}$;
    Update state path, $StatePath[t] \leftarrow S_t$;
    Update cost path, $CostPath[t] \leftarrow$ Cost(S_t, x_t, ω);
    Transition to new state, $S_{t+1} \leftarrow$ Transition(S_t, x_t, ω);
  for $t = \mathcal{T}, \ldots, 0$ **do**
    Get visited state, $S_t \leftarrow StatePath[t]$;
    Calculate value of state, $CumulativeCost \leftarrow \sum_{\tau=t}^{\mathcal{T}} CostPath[\tau]$;
    Update with true value, $\bar{V}_t^k(S_t) \leftarrow$
    UpdateValueApproximationFunction($S_t, CumulativeCost$);
Draw realization path,
$(W_1, W_2, \ldots, W_{\mathcal{T}+1}) \sim$ ExogenousProbabilityDistribution();
for $t = 0, \ldots, \mathcal{T}$ **do**
  Get decision, $x_t^* \leftarrow$ PolicyDecision(S_t);
  Observe exogenous information, $\omega \leftarrow W_{t+1}$;
  Transition to new state, $S_{t+1} \leftarrow$ Transition(S_t, x_t^*, ω);
return Decisions for each time period, $x^* = (x_t^*)_{t \in \mathcal{T}}$;

i.e. the initial state S_0 and the number of simulations to perform to update the VFA before determining the final solutions. Steps 1-14, constitute the learning of the VFA. First, the paths are initialized, which are later used to calculate the true value of being in a state when updating the VFA. During VFA learning, K planning horizons are simulated. After each planning horizon simulation, steps 5-10, the VFA is updated in steps 11-14. The exact update procedure is elaborated in Section 7.2.3. A decision for the upcoming time period is determined using the procedure TRAININGDECISION(S_t), The actual realization of the exogenous information W_{t+1} is observed, the paths are updated and the state is transitioned based on the decision and the observed information. Note that W_{t+1} is a realization during time period t, but first available for decision making in time period $t + 1$. Thus, the exogenous information realized during the last time period $W_{\mathcal{T}+1}$ is defined.

The final phase of Algorithm 1, steps 16-19, concerns solving the actual problem, where decisions are made using the POLICYDECISION(S_t) procedure. The distinction between decisions chosen during the VFA learning and in the final phase is made because of varying objectives. The goal during the final phase is to solve Equation 7.3, while during the VFA learning, the goal is to determine a satisfactory approximation of the true value function.

7.2.2 Epsilon-greedy learning strategy

To adequately approximate the value function, a sufficient number of different states must be visited. If too few are visited, the algorithm risks getting stuck in local optima. Because of the stochastic realizations of different bacteria dispersal rates, the value of being in a state is not necessarily the same in every realization. To ensure a good estimate for the value of a state, this state or similar states should be visited multiple times.

To ensure exploration, the policy procedure TRAININGDECISION(S_t) is introduced. It is not only incentivized to minimize the objective in Equation 7.3, but to explore new states. This is achieved by employing a three-phased *epsilon-greedy* strategy. During the first 10% of the iterations, the decisions are made randomly. The second phase consists of selecting a random decision with probability ϵ, a greedy decision with probability ϵ^2 or an ADP decision with probability $1-\epsilon-\epsilon^2$, i.e. employing the POLICYDECISION(S_t) procedure described in Section 7.3. The third and final phase consists of the last 5% of the iterations, in which only ADP decisions are made. By employing this strategy, the focus is initially on exploration. The second phase ensures randomized greedy policies more similar to the ADP policy than the purely random policy. This warrants local randomization, but are decisions still similar to that of the ADP policy. Lastly, the ADP policy is followed for several iterations to ensure good estimates of the states typically visited when employing the policy. A detailed description of TRAININGDECISION(S_t)

is found in Algorithm 7 in Appendix C.

The problem in Equation 7.3 is a mixed-integer nonlinear program (MINLP) when using the feasible region $\mathcal{X}_t$ defined in Section 6.2.3. Although the problem is solved for a single time period only, not for the entire horizon, the optimization problem is still computationally demanding. Selecting a random decision from the set of all feasible decisions uniformly would require computing every feasible solution. As the solution space is large, this would make the random decision procedure slow. To avoid enumerating every solution and still ensure sufficient exploration, a pseudo-uniform random solution generator is applied. The procedure uniformly draws the amount of resources to be allocated and then allocates resources to the regions in a random order. The amount allocated to each region is uniformly drawn from the remaining resources at the time. Detailed pseudocode for the procedure is available in Algorithm 8 in Appendix C.

7.2.3 Value function approximation implementation

The VFA is modeled as a neural network. The network aims to learn the cumulative future cost of being in a state, conditional on some policy being followed. Due to the curse of dimensionality, discussed in Section 4.3.1, a look-up table would be computationally impractical. The curse of dimensionality can be somewhat mitigated by using state aggregation. However, linear regression or neural networks allows for updates of the entire function, not just neighboring states, when using look-up tables. Because the exact functional form of the value function is not known, the neural network is employed as VFA representation, instead of linear regression with basis functions. With a neural network, the entire VFA is updated when the value of a decision taken in a given state is realized, not just the single state or state similar to it, as is the case of state aggregation. The value function approximation is updated using the procedure UPDATEVALUEAPPROXIMATIONFUNCTION$(\cdot)$. The procedure applies neural network backpropagation at the end of each planning horizon iteration. To reduce bias that would arise from estimating $\bar{V}_t^k(S_t)$ using the estimate $\bar{V}_{t+1}^{k-1}(S_{t+1})$, the target is the realized cumulative costs, not the value function estimates.

An issue with neural networks, particularly deep neural networks with a significant number of parameters, is the slow learning from each data point. To overcome this challenge, each data point from the most recent iterations are used as training data. Each update iteration, several minibatches are drawn, i.e. random traning points from the training data to learn from. Therefore, the neural network might train on the same data point several times, making it possible to learn sufficiently from each observation.

By adjusting the training data each iteration, the convergence should also be more efficient, only learning on data from the recent policy, not the exploratory policy from earlier iterations. To avoid overfitting of the last realization of dispersal rates, the training data

consists of the past three iterations, not only the last one.

The state used as input to the neural network is as presented in Section 6.2.2. However, because the demand for vaccination and disinfectant is the susceptible population and the demand for rehydration solution and antibiotics is the symptomatic infected population, the state will have redundant, replicate demand values. These are removed to ensure an efficient state representation and faster learning. In addition, to further increase the learning efficiency, the states are scaled to zero mean and unit variance, more formally:

$$x' = \frac{x - \bar{x}}{\sigma_x}, \tag{7.4}$$

where x' is the scaled feature, x is the original state value, $\bar{x}$ is the mean value and σ_x is the sample standard deviation. The standardization scaling avoids issues with large differences in feature range and ensures weighting each feature in the state equally in the beginning. The scaling is based on the training data, as described above.

7.3 Subproblem Solution Method

The following sections describe different approaches to solve the MINLP arising from Equation 7.3 at each time step t, thus resulting in different versions of the procedure POLICYDECISION(S_t). Section 7.3.1, Section 7.3.2 and Section 7.3.3 present three policies not utilizing the VFA. The solution approach presented in Section 7.3.4 is a local search heuristic employing the VFA to balance the need for immediate response with the need of resources for future responses, and may be regarded as the main POLICYDECISION(S_t).

7.3.1 Greedy demand approximation

When an epidemic outbreak has occurred and the number of infections and fatalities rises, one possible is to allocate as many resources as needed to decrease the number of infected and fatalities, and to try to contain the epidemic as quickly as possible. The greedy policy aims to do this by satisfying all the immediate demand. If this is not possible, the available resources are allocated based on infection ratios.

Algorithm 2 takes the current state and the kit size as inputs and outputs a feasible decision for time period t, with resources allocated greedily based on infection ratio. First, the immediate demand is observed. The kit size is accounted for and the demand is ceiled, i.e. rounded up, to ensure that it is satisfied. Further, the total demand for each intervention type and the infection ratio is calculated. The available resources are computed, and the resources are allocated based on the infection ratio. Because the resource kits are integer and the infection ratios are not, an allocation where regions receive the same

Algorithm 2: GREEDYPOLICYDECISION(S_t, I, q)

Input : State S_t,
Current number of symptomatic infected in each region, I,
Kit size q.

Output: Feasible decision x_t.

Observe immediate demand and calculate ceiled demand, $D \leftarrow \left\lceil \frac{S_t.demand}{q} \right\rceil$;
Calculate infection ratio for all regions and interventions, $d_i \leftarrow \frac{I_i(t)}{\sum_{j \in \mathcal{I}} I_j(t)}, \forall i \in \mathcal{I}$;
for *intervention type* $m \in \mathcal{M}$ **do**
    Calculate total demand, $D_m^{TOT} \leftarrow \sum_{i \in \mathcal{I}} D_{im}$;
    Calculate resources to be allocated, $R_m \leftarrow \min(D_m^{TOT}, S_t.resources[m])$;
    Allocate integer resources to based on infection ratios,
    $x_{tm} \leftarrow$ ROUNDING$((d_i)_{i \in \mathcal{I}}, R_m)$;
while *not* ISFEASIBLE(x_t) **do**
    Calculate marginal costs, $MC^-, MC^+ \leftarrow$ MARGINALCOST(x_t);
    Identify resource with least impact, $(i, m) \leftarrow \arg\min_{(i,m) \in \mathcal{I} \times \mathcal{M}}(MC^-)$;
    Reduce decision where it is most efficient, $x_{tim} \leftarrow x_{tim} - 1$;
return Decision in time period t, x_t;

resource ratio as infection ratio is not necessarily achievable. Therefore, a rounding procedure is employed. Although the demand is satisfied, capacity constraints may restrict the decision made, thus feasibility is checked. If the decision is infeasible, the expected marginal cost of removing one resource of each type in each region is calculated, and resource where the least lives are saved, i.e. the resource having the highest marginal cost, is removed.

The procedure MARGINALCOST(x_t), returns the marginal cost of decreasing and increasing one unit of resource of type m in region i, for all intervention types and resources. That is, $MC^- = (MC^-_{im})_{(i,m) \in \mathcal{I} \times \mathcal{M}}$, where MC^-_{im} is the expected difference in objective value by decreasing resource type m in allocated to region i by one unit. Similarly, $MC^+ = (MC^+_{im})_{(i,m) \in \mathcal{I} \times \mathcal{M}}$, measures the expected difference in objective value by increasing one unit resource of type m to region i. Note that for the greedy algorithm, the marginal cost is not calculated using the VFA, only the immediate cost defined by $C^E(\cdot)$. If the kit size is small, the number of resources removed per iteration in the while loop can be adjusted to take larger steps.

The procedure ISFEASIBLE(x_t), controls that if given a resource allocation x_t, does a personnel allocation and available facilities exist, such that the decision is feasible. The variable of interest is the resource allocation x_t, within the constraints of available medi-

cal facilities and personnel, being the supporting decision variables, y_t and z_t. Therefore, the actual allocation of personnel and establishment of medical facilities are beyond the scope of this thesis.

checks that given a resource allocation, x_t, does there exist a personnel allocation and available facilities, such that the decision is feasible. The interest is the resource allocation, within the constraints of available personnel and medical facilities, being the supporting decision variables, thus the actual allocation of personnel and establishment of medical facilities is not of interest in this thesis.

7.3.2 Naive horizon allocation approximation

In countries that have previously experienced cholera outbreaks, the epidemic model can be used as a probable approximation of the outbreak timeline without major adjustments of parameters as the outbreak unfolds. Thus, a medical resource allocation based on the forecast demand throughout the planning horizon can be used. The problem with this approach is that it neither takes into account the realizations of the dispersal rate nor the consequences of allocating the resources. Therefore, outbreaks late in the planning horizon may have been avoided if additional resources were allocated early on. The simulated outbreaks must also assume a strategy when simulating, and for the naive policy the strategy is to do nothing.

Algorithm 3: NAIVEPOLICYDECISION$(S_0, \hat{l})$

Input : Initial state S_0,
Assumed dispersal rate $\hat{l}$.

Output: Decisions for all time periods, $x = (x_t)_{t \in \mathcal{T}}$.

1 Set epidemic model decision-parameters to no intervention, $\nu, \beta, \phi, \theta \leftarrow (0, 1, 0, 0)$;
2 Set bacteria dispersal rate to desired value, $l \leftarrow \hat{l}$;
3 Project symptomatic infected at each week, $I \leftarrow$ SIMULATEEPIDEMIC(S_0);
4 Calculate infection ratio, $d_{ti} \leftarrow \frac{I_i(t)}{\sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} I_i(t)}$;
5 Allocate integer resources to based on infection ratios, $x_m \leftarrow$ ROUNDING(d, R_m)
return Decision for entire planning horizon, x;

Algorithm 3 takes the current state S_t and the desired dispersal rate $\hat{l}$ as input, and returns resource allocation decisions for all time periods and regions. First, it ensures no interventions are employed and assumes a dispersal rate in the epidemic model. Furthermore, it projects the infections at each week and calculates the infection ratio by simulating the epidemic using the assumed dispersal rate. Based on the infection ratios of each region for each time period, the resources are allocated using a rounding procedure.

7.3.3 Myopic allocation

The myopic allocation algorithm is inspired by the forecasting approach in Long et al. (2018). However, it is not the same allocation algorithm as their proposed myopic algorithm. In this thesis, the myopic policy projects the demand for the upcoming week, and allocates based on the projected demand, instead of the immediate, observed demand that GREEDYPOLICYDECISION() does. The benchmark policy ignores demand, personnel and facility feasibility to provide insights on performance when large number of resources are allocated early.

Algorithm 4: MYOPICPOLICYDECISION(S_t, q)

Input : State S_t,
Kit size q.
Output: Decision x_t.

1 Project demand and infections in upcoming week,
$\tilde{D}, \tilde{I} \leftarrow$ SIMULATEEPIDEMIC(S_t);
2 Calculate ceiled demand, $\tilde{D} \leftarrow \left\lceil \frac{\tilde{D}}{q} \right\rceil$;
3 Calculate infection ratio for all regions, $d_i \leftarrow \frac{\tilde{I}_i(t)}{\sum_{j \in \mathcal{I}} \tilde{I}_j(t)}, \forall i \in \mathcal{I}$;
4 **for** *intervention type* $m \in \mathcal{M}$ **do**
5 Calculate total demand, $\tilde{D}_m^{TOT} \leftarrow \sum_{i \in \mathcal{I}} \tilde{D}_{im}$;
6 Calculate resources to be allocated, $R_m \leftarrow \min(\tilde{D}_m^{TOT}, S_t.resources[m])$;
7 Allocate integer resources to based on infections ratios,
$x_{tm} \leftarrow$ ROUNDING$((\tilde{d}_i)_{i \in \mathcal{I}}, R_m)$;
8 **return** Decision in time period t, x_t;

Algorithm 4 take the current state S_t and a kit size q as input and returns a decision for time period t. The demand for the upcoming week is projected. Based on this demand, the resources are allocated in a similar manner as the greedy policy. Although the myopic policy projects future costs, and therefore looks further than the current state, it only does so for one week, hence the name *myopic*.

7.3.4 Local search heuristic

Using the ADP approach presented with Algorithm 1, Equation 7.3 is solved with a local search heuristic. The initial solution is based on the greedy strategy presented in Section 7.3.1. Algorithm 5 presents the procedure LOCALSEARCHPOLICYDECISION(S_t), used in step 17 in Algorithm 1.

Algorithm 5 takes the current state S_t, the kit size for bundling medical resources q and

Algorithm 5: LOCALSEARCHPOLICYDECISION(S_t, q, Δ)

Input : State S_t,
Kit size q,
Impact factor Δ.
Output: Best found, feasible decision x_t.

Get feasible initial decision, $x_t \leftarrow$ GREEDYPOLICYDECISION(S_t, q);
Calculate marginal costs of decreasing and increasing unit resource, $MC^-, MC^+ \leftarrow$ MARGINALCOST(x_t);
while $\min\{MC^-\} < -\Delta$ **do**
    Get index of lowest marginal cost, $(i, m) \leftarrow \arg\min\{MC^-\}$;
    Remove resource from index with lowest marginal cost, $x_{tim} \leftarrow x_{tim} - 1$;
    Calculate new marginal costs, $MC^-, MC^+ \leftarrow$ MARGINALCOST(x_t);
Get best resource transfer indices, $(i, j, m) \leftarrow \arg\min_{i,j,m}(MC^+_{im} + MC^-_{jm})$;
while $MC^+_{im} + MC^-_{jm} < -\Delta$ **do**
    Increase resource allocation, $x_{tim} \leftarrow x_{tim} + 1$;
    Decrease resource allocation, $x_{tjm} \leftarrow x_{tjm} - 1$;
    if ISFEASIBLE(x_t) **then**
        Calculate new marginal costs, $MC^-, MC^+ \leftarrow$ MARGINALCOST(x_t);
    else
        Prevent revisit to resource increase, $MC^+_{im} \leftarrow +\infty$;
        Revert to feasible allocation, $x_{tim} \leftarrow x_{tim} - 1$;
        Revert to feasible allocation, $x_{tjm} \leftarrow x_{tjm} + 1$;
    Get new best resource transfer indices, $(i, j, m) \leftarrow \arg\min_{i,j,m}(MC^+_{im} + MC^-_{jm})$;
return Decision in time period t, x_t;

an impact factor Δ, as input, and outputs a feasible decision for time t. When marginal costs are calculated in the local search, the VFA costs are included, in contrast to the previous algorithms. In this way, the future cost of decisions are accounted for in the decisions made at time t.

The first part of the algorithm, steps 3-6, performs marginal cost descent, reducing the resource of type m in region i as long as the resource is more useful, i.e. reduce more fatalities, in later time periods. The VFA will have some inaccuracy in its predictions, especially as the output can be thousands of fatalities. Thus, to reduce the risk of removing resources that have insignificant effect on the objective, the impact factor Δ is introduced, to ensure removal or transfer of resources that provides significant impact on the objective value. If the impact factor is set too high, the ADP policy will imitate the greedy policy, while if it is set too low, a significant number of redundant changes to the decision may occur.

In the steps 7-17, the regions and intervention type where a transfer of resources decreases the cholera-induced fatalities the most, are identified. When there exists a transfer that reduces the fatalities with a certain level Δ, the transfer of one unit resource of type m from region j to region i is performed, given that it is feasible. If it is not, the marginal cost of increasing resources of type m in region i is set such that the same indices will not be revisited. Feasibility is only checked for the transfer and not for the marginal cost descent, i.e. first part of the algorithm. This is because the initial decision is feasible and thus reducing resources allocated cannot induce infeasibility. The transfer, however, might require opening an additional facility in the region receiving more resources. Due to facility location constraints and personnel constraints, the solution might be infeasible, and therefore the feasibility must be evaluated.

7.4 Regional Decomposition

When the number of regions increases, the number of possible, resource allocation decisions increases. The search space becomes very large, resulting in a more difficult approximation of the value function and the epidemic simulations become more computationally demanding. To limit this, the problem is decomposed by region, where each region is solved as its own, distinct ADP problem, thus requiring its own VFA. This approach introduces problems regarding the bacteria dispersal. Since the regions are solved independently in parallel, the dispersal between them are not calculated and must be provided in advance, reducing the accuracy of the epidemic model.

Algorithm 6 presents an alternative solution method, replacing Algorithm 1. The algorithm uses the same ADP framework and procedures defined in previous sections, but to increase readability, details regarding the value function updates are omitted. In contrast

to the stage decomposition approach, in Algorithm 6, the resources are allocated across regions at the beginning of the planning horizon. With a smaller search space, the VFA should be easier to learn for each region, albeit at the expense of the need for an initial resource allocation procedure, reducing response flexibility. After the initial allocation, the problem can be solved in parallel using a similar approach as in Algorithm 1.

The initial allocation procedures should leverage as much information as possible. In the first iteration, there is no available information guiding the allocation. Therefore, the outbreak is simulated without interventions and resources are allocated based on the number of fatalities. In all consecutive iterations, the allocation is based on marginal benefit computed based on results from previous iterations.

The procedure REALLOCATION() transfers resources between the regions where the marginal benefit is the highest and lowest, in terms of resources per fatality. The rationale of this procedure is that resources should be spent where they save the most lives. Assuming diminishing marginal benefit, the reallocation should converge towards a similar marginal benefit level across all regions. For detailed pseudocode describing the REALLOCATION() procedure, see Algorithm 9 in Appendix C.

Algorithm 6: REGIONALDECOMPOSITION$(\mathcal{T}, R, \mathcal{I}, K, H)$

Input : Planning horizon $\mathcal{T}$,
Initial resource pool R,
Regions $\mathcal{I}$,
Value function update iterations K,
Reallocation iterations H.

Output: Best found decisions x.

for *each reallocation iteration* $h = 1, \ldots, H$ **do**
  if $h = 1$ **then**
    Calculate fatalities in each region i without interventions;
    $r^h \leftarrow$ Allocate resources R to each region i based on fatalities;
  else
    $r^h \leftarrow$ REALLOCATION$(R, \mathcal{I}, r^{h-1}, S)$;
  Initialize all states, $S_{0,i}.resources \leftarrow r_i^h, \forall i \in \mathcal{I}$;
  Set bacteria dispersal in all states based on bacteria concentration in previous iteration or 0 if first iteration;
  for *each update iteration* $k = 1, \ldots, K$ **do**
    Draw realization path,
    $(W_1, W_2, \ldots, W_{\mathcal{T}+1}) \sim$ EXOGENOUSPROBABILITYDISTRIBUTION();
    for *each region* $i \in \mathcal{I}$ **do**
      for *each time period* $t = 0, \ldots, \mathcal{T}$ **do**
        Get decision, $x_{ti} \leftarrow$ TRAININGDECISION(S_{ti});
        Observe exogenous information from realization path, $\omega \leftarrow \mathrm{W}_{t+1}$;
        Transition to new state, $S_{t+1,i} \leftarrow$ TRANSITION(S_{ti}, x_{ti}, ω);
      Update value function,
      $\bar{V}_{ti}^k \leftarrow$ UPDATEVALUEFUNCTIONAPPROXIMATION$(S_{ti}, \bar{V}_{ti}^{k-1})$
Draw realization path,
$(W_1, W_2, \ldots, W_{\mathcal{T}+1}) \sim$ EXOGENOUSPROBABILITYDISTRIBUTION();
for *each region* $i \in \mathcal{I}$ **do**
  for *each time period* $t = 0, \ldots, \mathcal{T}$ **do**
    Get decision, $x_{ti}^* \leftarrow$ POLICYDECISION(S_{ti});
    Observe exogenous information from realization path, $\omega \leftarrow \mathrm{W}_{t+1}$;
    Transition to new state, $S_{t+1,i} \leftarrow$ TRANSITION$(S_{ti}, x_{ti}^*, \omega)$;
return $x^* = (x_{ti}^*)_{(t,i) \in \mathcal{T} \times \mathcal{I}}$;

Chapter 8

Case Data

This chapter provides an overview of the 2010 Haiti outbreak, being the case studied in this thesis. Section 8.1 presents the background and data regarding the operational capabilities in Haiti. Section 8.2 presents the parameter values used in the epidemic model, determined from previous literature or otherwise estimated. The purpose of the case study is not to compare the strategies from the solution methods presented in this thesis with the strategies used during the actual outbreak. For that, the data on the exact containment strategies in each case is too scarce. Rather, it is to apply the solution methods in this thesis on a realistic case, thus providing both a decision-support tool and evaluation tool for future epidemic outbreaks, where the data on containment strategy is available. The data is collected from different sources, mainly major health and humanitarian organizations, such as World Health Organization (WHO), United Nations Office for the Coordination of Humanitarian Affairs (OCHA) and Medecins Sans Frontieres (MSF), as well as government websites.

8.1 Haiti Cholera Outbreak in 2010

In January 2010 a devastating earthquake hit Haiti with magnitude 7.0 on the Richter scale. Only months later, in October 2010, a cholera epidemic broke out in the small Caribbean country still recovering from the earthquake. The epidemic would come to cost almost 10 000 lives and affect of 820 000 people (Pan American Health Organization, 2020a). The suspected cholera cases started on October 16 along the upper Artibonite River. October 20, the Haitian government confirmed the cases and declared the cholera outbreak, an epidemic (United Nations, 2011). The initial response focused on the Artibonite and Centre departments of Haiti. Although the cholera bacteria starting the

outbreak in Haiti originiated from a distant geographic source and was a result of human activity (Chin et al., 2011), there are today still cases occurring, making the outbreak relevant when developing response policies for future outbreaks.

8.1.1 Geographic and demographic data

Haiti consists of 10 departments, i.e. first-level administrative regions. The population in each department is based on Institut Haitien e Statistique et d'Informatique (2015), a government report from 2015, but corrected for assuming a 1.67% yearly population growth, reported in Pan American Health Organization (2017). The distance between each department is calculated based on the coordinates of each department capital, similar to the approach in Long et al. (2018). A map of Haiti and its departments is shown in Figure 8.1.

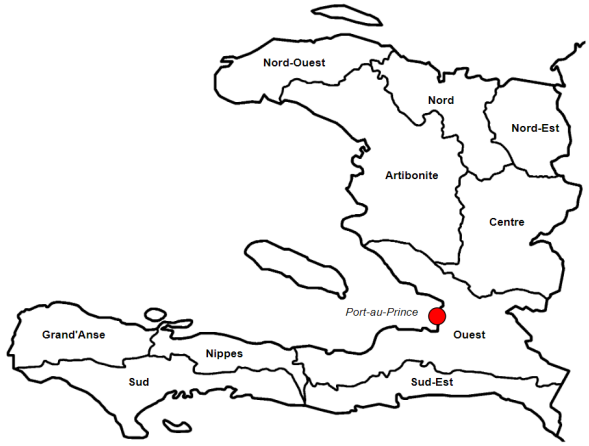

Figure 8.1: Departments of Haiti, also referred to as *regions* in this thesis. The capital, Port-au-Prince, is located in the department of Ouest.

8.1.2 Number of infections

Although previously available and reported in e.g. Bertuzzo et al. (2011), the government reports on number of cholera infections in 2010 is no longer available at the Haiti Ministry of Public Health (MSPP, 2020). The Pan American Health Organization (PAHO), the regional WHO in the Americas, released weekly situation reports during the first months of the cholera outbreak. The cumulative cholera cases are reported in cases per 10 000 inhabitants, at a national level and for the most affected departments (Pan American Health Organization, 2020b).

8.1.3 Treatment facilities

Accessed through the Humanitarian Data Exchange, the United Nations Office for the Coordination of Humanitarian Affairs (OCHA) shared the number of health facilities in Haiti, differentiating between hospitals, medical centers and dispensaries, as well as ownership and purpose, such as public, private and non-profit (OCHA Haiti, 2019). The type of facility is only available for the Ouest, Artibonite, Nord and Centre departments, but the total number of facilities in each department is reported nonetheless. To estimate the number of available facilities of each type in each department, all public and non-profit health facilities are assumed to be available for cholera treatment facilities during an epidemic. The ratio between for-profit, and public and non-profit facilities is calculated for the reported departments, and used to estimate the available facilities in the remaining departments. To estimate the number of CTCs, CTUs and ORPs available in the remaining regions, the ratio between hospitals, medical centers and dispensaries is calculated for the regions where reported. The mean ratios are used to estimate the number of CTCs, CTUs and ORPs in the remaining regions.

The treatment capacities are weekly capacity to treat or distribute resources. The CTC and CTU is assumed to have a treatment capacity of 100 and 30, respectively, based on the available beds described in Section 2.4.8. Therefore, the capacity for treatment resources, ORS and antibiotics, are assumed to be 100 and 30 for the CTCs and CTUs. ORPs are assumed to not treat with antibiotics, but have a buffer ORS treatment bed with capacity for two, for treatment until the patients can be transferred to a CTC or CTU. Distribution of vaccines and disinfectants is assumed to be the same for all facility types, being 1000. The unit time spent on distributing vaccines and disinfectants per allocated personnel is assumed to be 0.001, corresponding to spending 10 minutes on distributing a single resource. For ORS and antibiotics the unit time is assumed to be 0.07, corresponding to around 12 hours, because patients are hospitalized in a facility, requiring supervision, although not constantly.

The number of personnel required to operate a CTC, CTU and ORP is assumed to be 76, 16 and 2, respectively, based on the requirements described in Section 2.4.8.

8.1.4 Medical personnel

The PAHO emphasizes that Haiti lacks reliable data for personnel in health-related professions. In 2016, the Haitian Ministry of Heath assessed the available health professionals. According to the report, there are 15 980 health professionals in the public sector and 7 364 in the private sector (Pan American Health Organization, 2017). Assuming all medical personnel are mobilized during a national health crisis, this results in 23 344 available medical personnel.

8.1.5 Medical resources

According to OCHA, the medical resources available for containing the outbreak were tetracycline, an antibiotic, for 100 000 cholera treatments, rehydration salt for 200 000 treatments (OCHA, 2010a). The next day, OCHA also reported that 50 million chlorine tablets were available for water purification (OCHA, 2010b). Assuming each tablet purifies one liter of water and each person requires 20 liters of clean water per day, as described in Section 2.4.7, the chlorine disinfectant is sufficient for 2.5 million daily treatments, which corresponds to around 360 000 weekly treatments.

Early during the epidemic, the World Health Organization advised against oral cholera vaccines (OCV) as an emergency response (Pan American Health Organization, 2010a). Instead, they suggested mobilizing the response with rehydration treatment and improve water conditions. If vaccination were to be used, the WHO preferred a pre-emptive strategy over a reactive strategy, that is, they preferred vaccinating high-risk regions yet to have outbreaks instead of vaccinating susceptible people in the regions with active outbreaks. During the Haiti outbreak, the OCV stockpile had yet to be established, which can explain the reluctance to any vaccination campaigns. However, after the stockpile was established, 400 000 doses of OCV were allocated to Haiti in 2014 (Pan American Health Organization, 2014). A similar number is assumed to be available, were a new outbreak to occur.

8.1.6 Planning horizon

The planning horizon is assumed to be 120 days with weekly decisions, in accordance with the practice described in Section 2.3 and Section 2.4.3. The first decision is taken when the epidemic is official, and is thus made at $t = 0$. With weekly decisions, the number of stages is 18.

8.1.7 Medical kit size

In practice, medical resources will be sent in kits, not individual units. The kit size is the number of resources packaged together and is assumed to be 1000 units per kit for all resources. Further, it is assumed that kits consist of a single resources, e.g. vaccines and antibiotics are not placed together in the same kit.

8.1.8 Bacteria dispersal distribution

Bertuzzo et al. (2011) calibrated its epidemic model to the 2010 Haiti outbreak and found the bacteria dispersal rate to be $l = 0.025$. This is assumed to be the mode dispersal distribution. The two other scenarios is high rainfall, increasing the dispersal rate to $l = 0.25$ or drought, decreasing the dispersal rate to $l = 0.0$. Setting a reasonable distribution

is difficult without estimating it based on seasonal outbreaks, tying it to environmental factors such as weather.

8.2 Epidemic Parameters for the Haiti Case

Most parameters of the epidemic model can be determined from the literature. Cholera-specific parameters, such as the half-saturation constant, is set from a consensus observed in the literature. Case-specific parameters, such as the dispersal rate between regions, are set using the parameters from Bertuzzo et al. (2011), which studied the 2010 Haiti cholera outbreak and on which the network component of the epidemic model proposed in this thesis is based on. The values for the epidemic parameters used in the computational study are summarized in Table 8.1. For a more detailed explanation of each variable, see the model definition in Section 6.1.

The rate of vaccination, proportion of individuals consuming contaminated water, proportion receiving rehydration treatment and proportion receiving antibiotics are all parameters that depends on the decision from the resource allocation model, and thus do not have a given value. If no resources are allocated, the rate of vaccination and proportions receiving rehydration solutions and antibiotics would be 0, while the proportion ingesting contaminated water would be 1, since no disinfectant is allocated.

The water reservoir size is estimated in the same way as in Andrews and Basu (2011), except for transforming the metric to milliliters. The relative rate of cholera-induced death, when not receiving rehydration treatment is estimated based on information from World Health Organization (2018). Considering that the fatality rate of cholera in some untreated outbreaks is up to 50% and that symptomatic cholera can prove fatal within 12 hours, the cholera-induced death rate is calculated to 1.0 per day. Assuming the cholera-induced death rate when being treated with rehydration solution is $4.0 \cdot 10^{-3}$ per day, as in Bertuzzo et al. (2011), the increased fatality rate factor when not receiving rehydration treatment is assumed to be 250.

Table 8.1: Parameter definitions, values and references.

Parameter	Definition	Value	Reference
ν_i	Rate of vaccination	Decision-dependent	-
β_i	Proportion of individuals consuming contaminated water	Decision-dependent	-
ϕ_i	Proportion of individuals receiving rehydration treatment	Decision-dependent	-
θ_i	Proportion of individuals receiving antibiotics	Decision-dependent	-
W_i	Water reservoir size	$15 \cdot N_i \cdot 365 \cdot 10^3$ ml	Andrews and Basu (2011)
τ	Vaccine efficacy	0.82	World Health Organization (2017b)
κ	Half-saturation constant	10^6 cells/ml	Lemos-Paião et al. (2017)
p	Proportion of infected being asymptomatic	0.80	World Health Organization (2018)
γ	Rate of recovery	0.2 day^{-1}	Hartley et al. (2005)
ρ_A	Rate of excretion, asymptomatic individuals	$1.3 \cdot 10^8$ cells/day	Andrews and Basu (2011)
ρ_I	Rate of excretion, symptomatic individuals	$1.3 \cdot 10^{11}$ cells/day	Andrews and Basu (2011), Kaper et al. (1995)
χ	Relative rate of cholera-induced death, not receiving rehydration treatment	250	Estimated from World Health Organization (2018)
λ	Relative rate of recovery, receiving antibiotics	2.3	Andrews and Basu (2011)
ψ	Relative rate of excretion, receiving antibiotics	0.52	Andrews and Basu (2011)
μ	Rate of non-cholera induced death	$4.6 \cdot 10^{-5}$ day^{-1}	Bertuzzo et al. (2011)
μ_B	Rate of cholera bacteria decay	0.03 day^{-1}	Andrews and Basu (2011)
μ_C	Rate of cholera-induced death	$4.0 \cdot 10^{-3}$ day^{-1}	Bertuzzo et al. (2011)
D	Mean dispersal distance	9.0 km	Bertuzzo et al. (2011)
l	Rate of regional bacteria dispersal	$2.5 \cdot 10^{-2}$ day^{-1}	Bertuzzo et al. (2011)

Chapter 9

Computational Study

The purpose of this thesis is to develop a decision-support tool to efficiently respond to epidemic outbreaks. This chapter investigates to what extent applying the solution methods described in Chapter 7 fulfills the thesis' purpose. Section 9.1 presents the calibration of the epidemic parameters to fit the Haiti case presented in Chapter 8. Next, Section 9.2 presents the hyperparameter tuning of the value function approximation (VFA) and the convergence results.. The efficiency of different medical resource allocation policies applied to the calibrated epidemic model is investigated in Section 9.3. An alternative epidemic outbreak is constructed and analyzed in Section 9.4. Lastly, Section 9.5 investigates the impact on policy selection when varying key parameters in the epidemic and resource allocation models.

The solution methods are implemented in Python 3.7.6. The data is handled using Numpy 1.18.2 and Pandas 1.0.3, the differential equations defined in Section 6.1 are solved using Scipy 1.4.1, feasibility is checked using OR-tools 7.7 and the neural network used for VFA is implemented with PyTorch 1.4.0. All computations are performed on an Intel® Core™ i7-8700T 2.40GHz CPU with 16GB RAM. An overview of the code structure is presented in Appendix A.

Chapter 7 proposes two different approximate dynamic programming (ADP) approaches, one with a single, complex VFA, solving the problem combined for all regions at each stage, and the other approach allocates resources to regions at the beginning of the outbreak, letting each region allocate their resources over the planning horizon. In the latter approach, there is one VFA for each region, allowing for parallelization, at the cost of flexibility, not allowing for reallocation between regions as the bacteria dispersal rate is realized. Throughout the computational study, the ADP policy refers to the policy using

the *stage decomposition* approach. After the methods were developed, a significant improvement in the computational performance of the epidemic model was implemented, by assuming constant population size when calculating dispersal probability between two regions, as described in Section 6.1. This improvement thereby removed the original motivation for the *regional decomposition* approach.

9.1 Epidemic Model Calibration

All suspected cholera cases should be treated, but a large-scale response to the outbreak occur shortly after an cholera outbreak is confirmed. The first available data on cumulative cases of infections for the 2010 Haiti outbreak is from October 28, 2010, more than a week after the Haitian government declared an epidemic on October 20. By then, a large-scale response should already have been begun, thus, it is insufficient to be the initial condition. On October 21, the Pan American Health Organization (PAHO) issued a press release on the cholera outbreak, reporting 1500 cases of symptomatic cholera in the department of Artibonite (Pan American Health Organization, 2010b). This is sufficiently close to the epidemic confirmation date for it to be a natural starting point for a large-scale response to the epidemic, and is thus set as the initial condition at time $t = 0$. Being the day after laboratory confirmation, any available vaccines are expected to arrive at time $t = 13$, according to the International Coordinating Group on Vaccine Provision (ICG) timeline described in Section 2.4.4.

In the literature, the initial concentration of cholera bacteria in the water of region i, $B_i(0)$, is typically calibrated to the case. However, the calibrated value is rarely reported. To get a realistic initial condition, $B_i(0)$ is estimated using the cumulative cholera cases, both at the national level and in the regions with reported data. How the bacteria were introduced in the first place and reached the level of $B_i(0)$ is not investigated further and considered outside the scope of this thesis. The subject of interest is the response, beginning when the cholera outbreak is official.

Note that the real cumulative cholera infections are a result of the actual response to the outbreak. The decision-dependent parameters in the epidemic model would therefore vary during the actual outbreak. Calibrating each decision-dependent parameter at every time period is very difficult, would be very computationally demanding. Therefore, a constant strategy is assumed when calibrating the initial condition to the data. As mentioned in Section 8.1.5, the WHO advised against vaccines, and as explained in Section 2.4.6, antibiotics is not applied as treatment except for the most severe cases, to avoid antibiotic resistance. Therefore, vaccines and antibiotic treatment are ignored when fitting the data. It is further assumed that 10% have access to completely uncontaminated water and oral rehydration solution (ORS) is accessible to all reported symptomatic infected.

The outbreak began in Artibonite and quickly spread to regions close to the Artibonite River, being Centre and Ouest, Nord and Nord-Ouest. The initial bacteria concentration for these regions are found using the Nelder-Mead simplex algorithm from Scipy 1.4.1. The quality of fit is measured using the sum of squared errors (SSE). The initial conditions for the non-zero cholera concentration regions is reported in Table 9.1.

Table 9.1: Initial conditions for epidemic model for non-zero cholera concentration regions.

Parameter	Artibonite	Centre	Nord	Nord-Ouest	Ouest
$S_i(0)$	1 590 230	686 929	982 364	670 886	3 709 447
$A_i(0)$	6000	0	0	0	0
$I_i(0)$	1500	0	0	0	0
$R_i(0)$	0	0	0	0	0
$M_i(0)$	138	0	0	0	0
$B_i(0)$	4873	3003	634	849	1911

With the parameters determined based on the literature and calibration, and assuming a constant bacteria dispersal rate, $l = 0.025$, as calibrated in Bertuzzo et al. (2011), the estimated epidemic evolution compared to collected data of cumulative reported symptomatic cases are shown in Figure 9.1. Observe that early reported cases are consistently below the estimated cases, and late cases are above.

A possible explanation for predicting too many cases early, is inadequate reporting at the beginning of the outbreak, resulting in many unrecorded cases. It is difficult to get an overview of every case, especially in rural areas, because certain cholera symptoms, such as diarrhea, are not uncommon and and plausibly caused by other phenomena and diseases, such as the norovirus or rotavirus. However, another explanation is that too few variables are described. The model fails to provide an initial exponential growth that is both typical at the beginning of outbreaks and observed in the data. A possible solution is to include more variables in the calibration, i.e. increase the degrees of freedom.

The results of the parameter calibration are restricted by the quality of the data the parameters are tuned on. The data quality is sufficient, but questionable. For instance, in the reported data, the cumulative cases in Centre decrease form one week to another, which should not be possible. This is adjusted for by discarding the first data point, under the assumption that recent data is more accurate due to better overview of the outbreak. The purpose of the epidemic model is not to perfectly predict the 2010 Haiti outbreak, but rather produce realistic outbreak scenarios that the VFA can learn from. Based on the findings in Figure 9.1, the parameters partially fulfills this purpose, i.e. are sufficiently reasonable and realistic.

The epidemic model aims to produce realistic outbreak scenarios for the resource allocation model, but not to perfectly predict the 2010 Haiti outbreak. Based on the findings in Figure

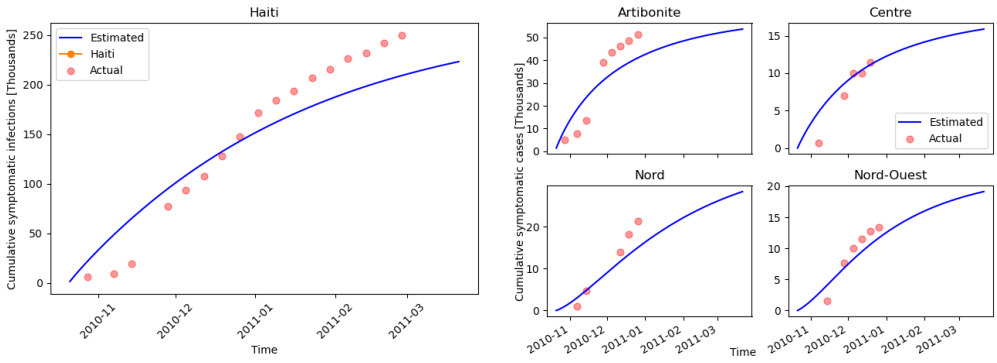

Figure 9.1: Estimated and actual cumulative symptomatic infections in Haiti and departments where data is available during the 2010 cholera outbreak.

To further support this claim, Figure 9.2 shows the predicted progression of the epidemic in the Haitian departments. In the beginning, the departments of Artibonite and Ouest have a significant increase that lasts throughout January 2011. The Ouest outbreak is large, but the region still has less absolute cumulative symptomatic cases than Artibonite, the latter of which which has a significantly smaller population. The simulated epidemic spreads rapidly, reaching eight of the ten departments in four days. The southwestern departments remain relatively isolated from the outbreak, but as is shown in Figure 9.3, the projected epidemic reach all regions by December 2, 2010. During the actual outbreak all regions were reached between November 14 and 21 (Pan American Health Organization, 2020b). Although the epidemic does not spread identically, they are similar, thus fulfilling the purpose of the epidemic model of realistic disease spread.

Note that it is difficult to make a realistic comparison with the actual response to the 2010 Haiti outbreak in terms of fatality, because the epidemic is not calibrated to fatalities, due to lack of data. Thus the parameters defining the transition rate from symptomatic to fatality, might be wrong. The more realistic both the number of infections and fatalities are, the better. However, as the primary purpose is to compare an ADP policy with other benchmark policies, not the actual outbreak response, the lack of calibration to actual fatalities is not considered a substantial issue.

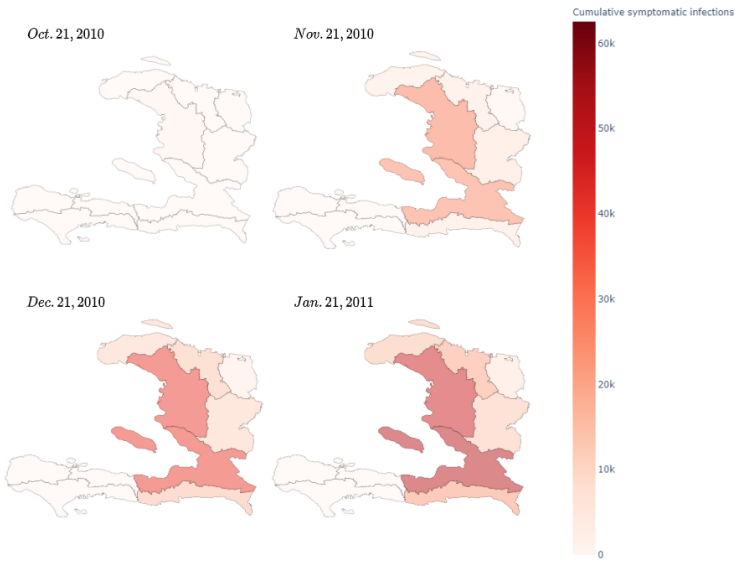

Figure 9.2: Projected cumulative symptomatic infections at various dates.

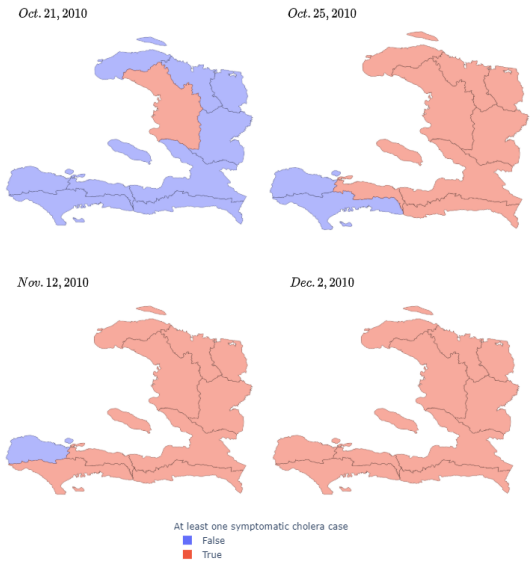

Figure 9.3: Departments with projected symptomatic infections at various dates.

9.2 Value Function Tuning and Convergence

The following subsections present the results from the hyperparameter tuning, reported in Section 9.2.1, and the results when investigating the ADP policy's convergence in Section 9.2.2.

9.2.1 Hyperparameter tuning

This section calibrates selected hyperparameters of the neural network used to represent the value function approximation (VFA). The hyperparameters tuned are the learning rate and the network architecture, in terms of the number of hidden layers and the neurons in each layer. The other hyperparameters and modeling choices are set using default best-practice values or empirical trials and are reported in Table 9.2. The number of minibatches drawn per update, i.e. the epochs, is set to 10. Although this is a low value, a too high value could overfit the VFA, due to the training data only being values of recently visited states, as explained in Section 7.2.3. If the VFA is having trouble converging, increasing the number of update iterations K, is a more robust approach.

Table 9.2: Hyperparameters

Hyperparameter	Value
Activation function	ReLU
Loss function	MSE
Optimization function	Adam
Output function	Linear
Minibatch size	32
Epochs	10
Impact factor Δ	500
VFA update iterations K	200

The results from the hyperparameter tuning are summarized in Table 9.3. The tuning took about 74 000 seconds, thus the average training time of the VFA is 7 400 seconds, slightly more than 2 hours.

While the training loss may indicate performance, it is the out-of-sample test loss that is of interest when evaluating the generalization performance. For the learning rate, the learning rate of 0.1 provides the lowest test loss. A learning rate of 1.0 provides severely worse performance, which can be explained by oscillation due to a too large step size. The loss minimum may be steep, thus the too large learning rate will step over it. An unexpected result is for a learning rate of 0.01, where the loss is significantly higher than for 0.001 and 0.1. A significantly higher test loss compared to training loss would typically indicates overfit. Because the training loss is on the same scale as for other learning rates, while the test loss is significantly higher than for similar learning rates, overfiting does not appear to be the case. The policy changes over the learning iterations, thus the VFA

Table 9.3: Mean squared error (MSE) loss in millions for various hyperparameters on training and test data.

Hyperparameter	Value	Loss	
		Training	*Test*
Learning rate	0.0001	2.82	3.93
	0.001	1.77	3.56
	0.01	1.75	26.7
	0.1	2.14	1.69
	1.0	3.25	111
Architecture	[]	2.06	1.98
	[10]	2.81	2.07
	[20, 10]	1.82	2.53
	[30, 20, 10]	3.72	1.58
	[40, 30, 20, 10]	2.42	2.20

could become stuck for some states, predicting values based on previous realizations, but not what would occur with the current policy. To avoid this, the number of learning iterations could be increased above $K = 200$ when investigating the convergence.

The architecture test loss is similar for various architectures. The minimum loss appears to be for architectures with three hidden layers, with 30, 20 and 10 neurons in the layers respectively. The training error for the 30, 20, 10 architecture is higher than the test loss, thus overfitting does not appear to be an issue. In the following sections, the architecture used is [30, 20, 10]. The policy convergence is tested with a learning rate of both 0.001 and 0.1 and the number of update iterations are increased to $K = 500$.

9.2.2 Policy convergence

Using the hyperparameters found above, the convergence of the resulting ADP policy is investigated. The VFA is conditioned on some policy. As it learns it should provide more accurate predictions as to how many fatalities will occur if the policy is followed, thus the prediction loss should decrease. In addition, the predictions should decrease as the policy improves. Eventually, it is expected to converge, providing consistent predictions given a state with low prediction loss.

Figure 9.4 shows the predicted cost, i.e. the future cholera-induced fatalities, given the initial state, when using the ADP policy, as well as the loss of the prediction measured as the mean squared error (MSE). Thus, the VFA predicts the total cholera-induced fatalities throughout the planning horizon, when following the ADP policy. The number of learning iterations are $K = 500$, the learning rate is 0.001 and 0.1, and the other parameters are set as in the hyperparameter tuning.

For the first 50 iterations, the VFA is trained purely on random decisions, as explained in Section 7.2.2. As it learns, the prediction loss drastically decreases, until about 100

iterations for a learning rate of 0.001 and almost immediately for a learning rate of 0.1. After the first 50 iterations, the number of random decisions are still high, but decisions are made using the ADP and greedy policy as well. With a learning rate of 0.001, the loss remains low after the 100 iterations, and the predicted fatalities decreases, thus confirming that the ADP policy is in fact changing towards the better. Between iterations 200 and 500, the predictions are relatively stable, and in this interval more and more decisions are made using the ADP policy. In contrast, with a learning rate of 0.1, the predictions oscillate significantly during most of the iterations. The last 25 iterations are made purely with the ADP policy, which explains why the oscillations are reduced further for both learning rates, at the end. Observe that the loss remains small, and minor increases are further reduced as the iterations increase. Although less iterations are required to decrease the initial loss with learning rate 0.1, the learning rate appears too large to converge properly during most of the iterations. Both the investigated learning rates appear to converge towards around 20 000 predicted fatalities, hence there are no major performance differences. Therefore, learning rate 0.001 is used when training the models for the rest of the computational study.

Overall, the ADP policy appears to converge, especially with a learning rate of 0.001. It consistently predicts a little below 20 000 cholera-induced fatalities throughout the planning horizon and the prediction loss is small. However, the decisions made with the ADP policies at each stage are not necessarily optimal. Thus, even though the VFA converges to making consistent decisions, it does not necessarily mean the policy reached is optimal.

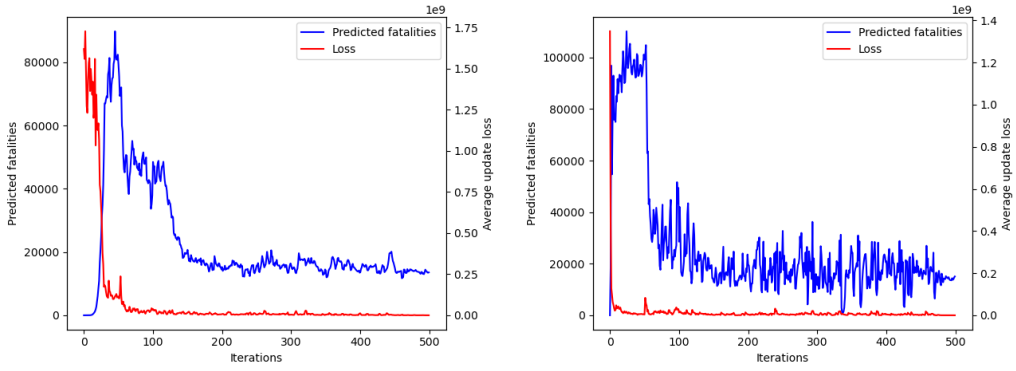

(a) Learning rate 0.001

(b) Learning rate 0.1

Figure 9.4: Value function approximation (VFA) fatalities prediction and average loss for the initial state of the 2010 Haiti outbreak, for each VFA update iteration, using a learning rate of (a) 0.001 and (b) 0.1.

9.3 Resource Allocation Policies Efficiency

In the following sections, unless otherwise specified, the models are applied to the case described in Chapter 8, henceforth referred to as the *base case*.

The exogenous information process, i.e. the stochastic bacteria dispersal rate, ensures different realized epidemic outbreaks. Therefore, it is not sufficient to conclude on performance based on a single realization. Instead, the policies are simulated multiple times to approximate expected performance. The mean performance for the ADP, greedy, naive and myopic approaches across 100 simulations are shown in Figure 9.5. The 95th and 5th percentiles are also included. In Figure 9.6, histograms of the fatalities under the different policies are included. The realization paths explored are the same for all policies in the histograms. Note that the impact factor used is decreased to $\Delta = 20$, due to less time constraints. The value is based on empirical trials, balancing the trade-off between employing the VFA in the local search and making redundant changes to the decision due to minor inaccuracies in the VFA.

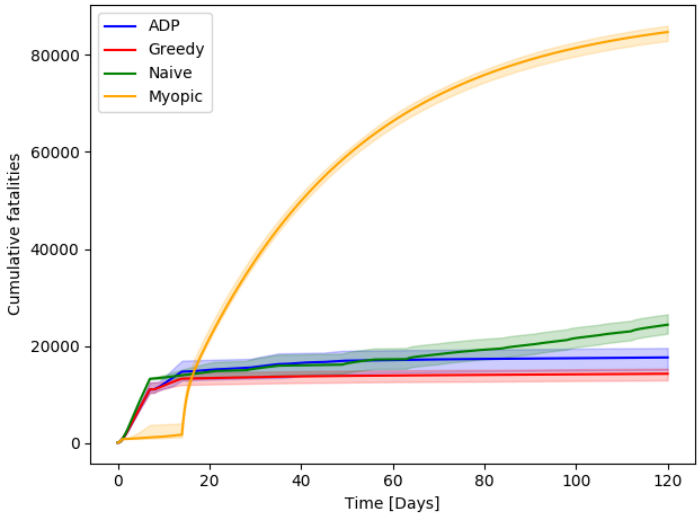

Figure 9.5: Mean cumulative fatalities under the ADP, greedy, naive and myopic resource allocation policies, across 100 epidemic realizations. The interval edges are the 95th and 5th percentiles performance of the respective policy.

Naive policy: The naive policy simulates the epidemic with a given dispersal rate when no interventions are made, and allocate resources based on the demand in a specific region at a specific point in time over the planning horizon.

When there is little resource scarcity, the naive policy is expected to perform well. In

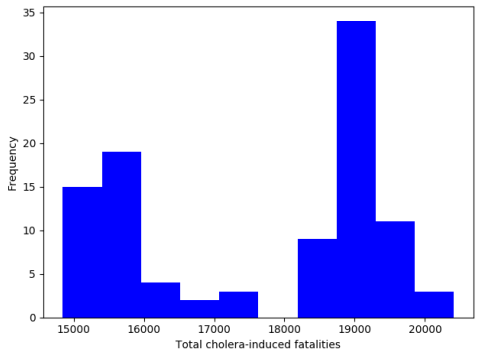

(a) ADP policy

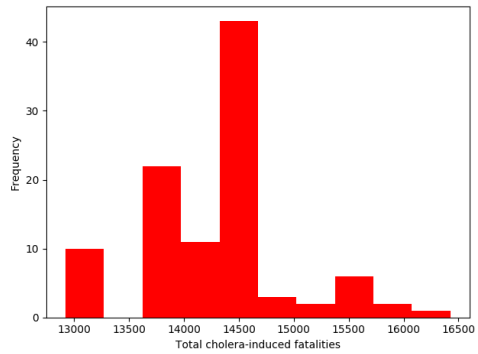

(b) Greedy policy

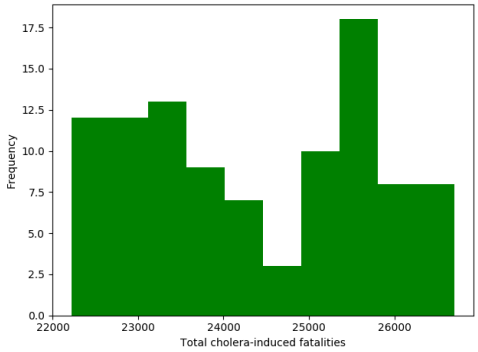

(c) Naive policy

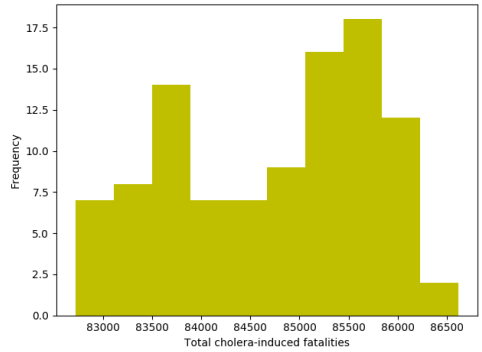

(d) Myopic policy

Figure 9.6: Frequency of cholera-induced fatalities for different bacteria dispersal realization path when employing the different policies, for 100 simulations.

that case, there are enough resources to satisfy all demand, thus ensuring both containment and treatment of later infections. However, as the scarcity increases, focusing on containing the epidemic by allocating more resources early, might be a better strategy. More resources early can reduce the number of infections, and thus reduce the overall magnitude of the outbreak. This early containment strategy is captured by the greedy policy. From Figure 9.5, the naive policy appears to allocate less resources than the greedy. Similar to the ADP policy, the naive thus have worse performance during the first days. However, at the final days the fatalities with the naive keeps increasing, while it stays close to constant for the ADP policy, indicating that the ADP policy better identifies what resources to save for later use, compared to the naive policy. The histogram in Figure 9.6c supports this. Although the distribution looks bimodal, the mass is distributed more evenly than both the ADP and greedy policy. The naive policy determines

the allocation based on a simulation prior to the outbreak, and thus does not adjust to the outbreak. A lack of adjustment increases the sensitivity to stochasticity, which explains the more uniform mass distribution. Another explanation for the bimodal appearance of the naive policy is too few simulations. The 100 simulations may not be sufficient to provide a reliable mean performance.

Greedy policy: The majority of the fatalities under the greedy policy occurs in the first week. The initial outbreak started with only infected individuals in Artibonite, therefore, resources are only allocated to this region. However, during the first week, the epidemic spread to several other regions. Without oral rehydration solutions (ORS) allocated to these regions, the cholera-induced death rate is very high. Although the known symptomatic cases were in Artibonite, it is possible that there were already other infected individuals in other regions. With more information from other regions, the greedy policy would capture this and ensure a lot fewer fatalities. Starting the response at an earlier point in time, before the epidemic is announced, would give a better overview of the number of symptomatic infected in other regions. Thus, the greedy policy would have sufficient information to allocate more resources early on, and possibly perform better by avoiding the drastic increase during the first week, as seen in Figure 9.5.

In terms of robustness, the greedy policy ensures consistent results on the base case, as seen in Figure 9.6b. The mean is about 14 500 fatalities, and only a few realization paths leads to more fatalities. Due to focusing on containing the outbreak, instead of ensuring sufficient resources if the outbreak reemerged, the greedy seems to be succeeding in the containment. If that is the case, it also explains the robustness of the greedy policy, because of the outbreak is contained, then the later realizations of bacteria dispersal are negligible, leading to consistent performance in terms of fatalities.

Myopic policy: The myopic policy forecasts the demand for the upcoming week and allocates greedily based on symptomatic infections in each region. The initial fatalities are reduced drastically, as it projects what regions will receive infections throughout the week. Although it provides a more rapid response to the outbreak, and greatly reduces the number of fatalities early on, the resources are depleted very early, resulting in a surge in fatalities later on.

Due to its lack of resources later on, the myopic policy fails to adjust to any changes in bacteria dispersal. This is seen in Figure 9.6d, where the mass distribution is more uniform compared to the ADP and greedy policy. Due to its low performance and thus relevance, and to increase the readability of future plots, the myopic policy is omitted wherever deemed necessary. However, its performances are included in tables for reference and in plots in Appendix D.

ADP policy: For the base case, the ADP policy performs consistently worse than the greedy policy. A possible explanation for this, is that there are no reemerging outbreaks

throughout the planning horizon. Thus, after making similar decisions in the beginning, the resources saved for later usage are rarely utilized. Were outbreaks to reemerge, the performance would possibly improve, because the saved up resources can be allocated to treat a higher proportion of the infected. However, if the greedy policy is in fact close to optimal, then, ideally, the ADP policy would capture this and converge towards the greedy policy, instead of saving resources for later usage.

The distribution of fatalities given different bacteria dispersal realization paths seem bimodal, as seen in Figure 9.6a. This is unexpected and could be due to too few realizations. Due to time constraints, more realizations are not run. The current number of iterations does not seem to be enough to draw conclusions concerning the true distribution of fatalities under various policies, but it does seem to be sufficient to conclude that the greedy policy performs consistently better than the ADP policy, which in turn performs better than the naive and myopic policies. The greedy mean is 14 500, the range of ADP is 15 000-20 000, the naive policy range is 22 000-27 000, and the myopic is 83 000-87 000.

Similar to the greedy policy, most of the fatalities under the ADP policy occur during the first week. This is likely due to the lack of information on initial cases in other regions. Since the initial number of symptomatic infections in Artibonite is already high, and the bacteria concentration in the water reservoirs are significant, it is likely that there are other cases in other regions already. The model does not capture this during the first decision, due to the demand constraint. Delaying the response a single day allows the model to estimate the likely number of infected people in other regions, ensuring a more efficient response. This is confirmed in Figure 9.7, which shows a drastic decrease in fatalities under the ADP and greedy policies, compared to the immediate response in Figure 9.5. Instead of delaying the response and assuming zero infections in all regions except Artibonite, because there was no suspected cases there at the time, the initial number of symptomatic and asymptomatic infections in the regions could be included as variables in the calibration of the epidemic model. Observe that although the naive also performs better at the beginning in Figure 9.7, it appears to not sufficiently contain the epidemic, causing an increase in fatalities later on.

9.4 Alternative Epidemic Outbreak

For the Haiti case, the resource allocation model starts the response at the time an outbreak is occurring. The following section assumes a different epidemic setting. The same operational conditions as in the base case are assumed, but a higher bacteria excretion rate for both asymptomatic and symptomatic infected individuals is used, that is, $\rho_A = 7.3 \cdot 10^8$ and $\rho_I = 7.3 \cdot 10^{11}$, respectively. The increase from $\rho_A = 1.3 \cdot 10^8$ and $\rho_I = 1.3 \cdot 10^{11}$ is large, but still within the range from 10^{11} to 10^{13}, identified for symptomatic infected individuals with acute cholera in Kaper et al. (1995). This results

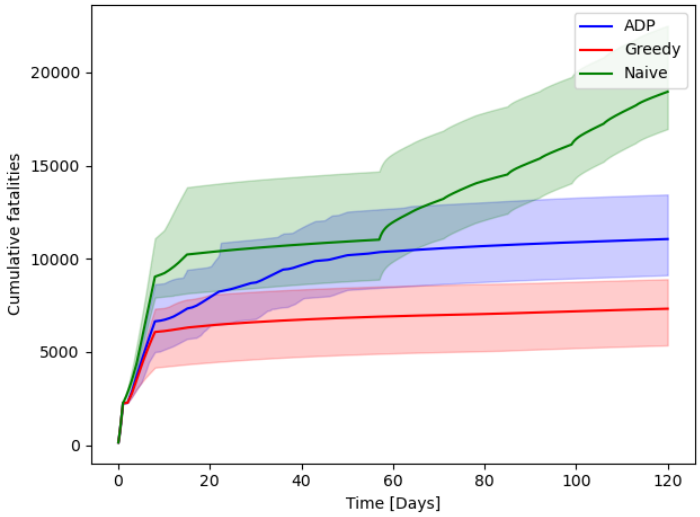

Figure 9.7: Mean cumulative fatalities under the ADP, greedy and naive policies. The interval edges are the 95th and 5th percentiles performance of the respective policy for 100 simulations.

in significantly larger and more unstable outbreaks, because each infected person excretes more bacteria, which in turn increases the likelihood of other individuals getting infected. The cumulative cases from the first decision $t = 0$ under the same assumptions as in Section 9.1, are shown in Figure 9.8. The mean cumulative fatalities, as well as the 95th percentile, across 100 iterations for the different policies are shown in Figure 9.9. In the following sections the VFA is trained using the same hyperparameters as the base case, except for an increase in the impact factor to $\Delta = 1000$, to adjust for the increase in realized costs. The increase causes larger inaccuracy in the VFA predictions, thus a too low value will make redundant changes to the decision. The value is again set using empirical trials, balancing the trade-off between employing the VFA in the local search and making redundant changes to the decision due to inaccuracies in the VFA.

The alternative epidemic outbreak is more sensitive to realizations of the bacteria dispersal rate, compared to the base case analyzed in Section 9.3. In particular, the ADP, greedy and naive policies have larger intervals for the 95th and 5th percentiles at the end of the planning horizon, compared to Figure 9.5. Investigating the frequency of fatalities for the different dispersal realization paths, the ADP policy seems more robust than the greedy policy. Figure 9.10. By rationing resources for later utilization, the ADP get consistent results around 90 000 fatalities. In a few realizations, the reemerging outbreak is very large and the ADP policy depletes its resources too early and the fatalities increase drastically. The greedy policy have a more uniform mass distribution, suggesting that it adjusts worse to the bacteria dispersal realization. The realization paths are the same for

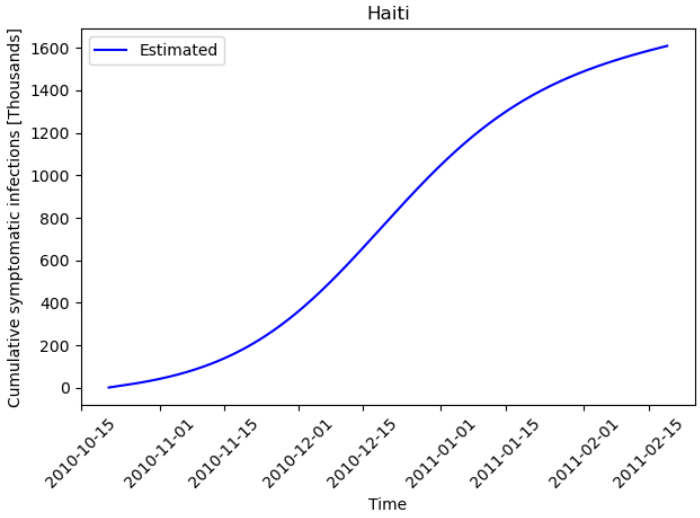

Figure 9.8: Cumulative cases in alternative epidemic outbreak with same assumptions as in base case.

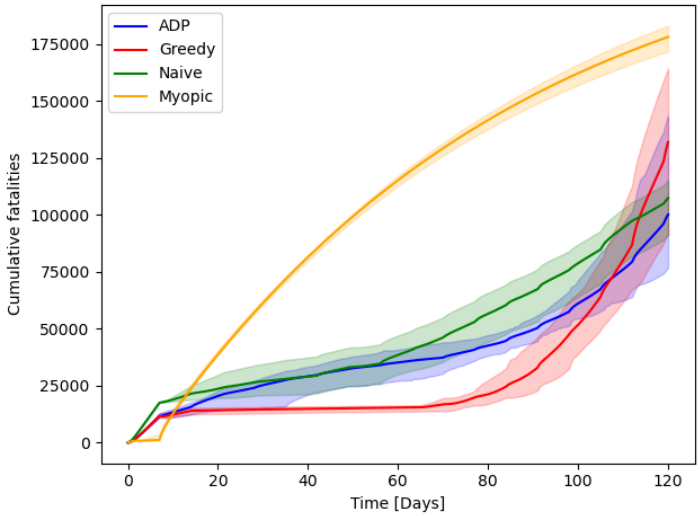

Figure 9.9: Comparison of policies with mean cumulative fatalities across 100 epidemic realizations. Intervals are 95th and 5th percentiles of the respective policy.

the ADP and the greedy policies, thus it is clear that there are several realization paths that cause drastic increases in fatalities for the greedy policy, that the ADP policy robustly responds to.

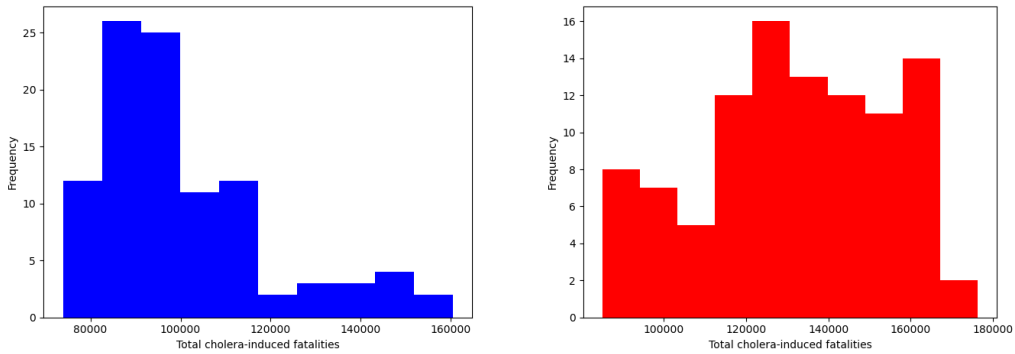

(a) ADP policy

(b) Greedy policy

Figure 9.10: Frequency of cholera-induced fatalities for different bacteria dispersal realization path when employing the different policies for 100 simulations.

Although the greedy policy performs best for almost the entire planning horizon, the mean cumulative fatalities are fewer for the ADP approach at the end. As the excretion rate of symptomatic and asymptomatic infections are higher, there is an increased probability for reemergence of the outbreak.

The allocation of medical intervention resources aggregated for all the regions, for the ADP and greedy policy, are shown in Figure 9.12. The allocation of disinfectants and vaccines are close to identical. The ADP policy allocates some vaccines later. The most notable difference is the allocation of ORS resources. The ADP policy allocates significantly less than the greedy policy for the third week. The following weeks, the ADP allocations are also noticeably less. The ADP policy also allocates less at the beginning of the reemerging outbreak in week 13, and instead wait until week 16 to allocate the most ORS. In this specific realization of the epidemic, the ADP policy decreased the total fatalities with as much as 29%, from about 129 000 to about 97 000 fatalities.

Table 9.4 shows the allocation of rehydration and antibiotics employing the ADP and greedy policy in weeks 2, 13 and 16. Having established from Figure 9.11a that the ADP reduces the allocation of ORS early to allocate later, note that the ADP policy, in contrast to the greedy policy, does not allocate as many resources to Ouest initially. Even when the epidemic is reemerging in week 13, the ADP policy holds off a large allocation, ensuring it has enough rehydration solutions to allocate throughout the last week of the epidemic as well. Ouest has the highest population of all the regions in Haiti. A possible explanation as to why the ADP policy keep the allocations to Artibonite, but reduces the resources allocated to Ouest, is that the percentage infected of the regional population is higher in Artibonite than Ouest. Thus, if the resources are scarce, it might be preferable to allocate based on percentage infected instead of absolute infected in the region.

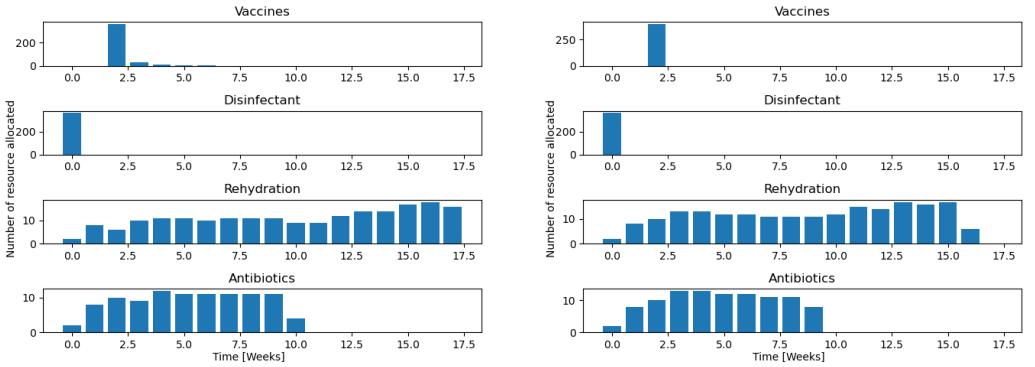

(a) ADP policy

(b) Greedy policy

Figure 9.11: Aggregated allocation of intervention resources across regions in one realization path for (a) the ADP policy and (b) the greedy policy.

Table 9.4: Rehydration solutions allocated for week 2, 13 and 16 for different policies (ADP / Greedy) for one realization path.

Regions	Week 2	Week 13	Week 16
Artibonite	2 / 2	2 / 3	3 / 1
Centre	1 / 1	1 / 1	1 / 1
Grande Anse	0 / 0	1 / 0	1 / 0
Nippes	1 / 1	1 / 0	1 / 1
Nord	0 / 1	2 / 4	0 / 1
Nord-Est	1 / 1	0 / 1	1 / 0
Nord-Ouest	1 / 1	2 / 0	2 / 0
Ouest	0 / 2	3 / 5	6 / 1
Sud	0 / 0	1 / 0	0 / 0
Sud-Est	0 / 1	1 / 3	3 / 1

Overall, on the alternative epidemic outbreak, the ADP policy provides a robust policy with the least fatalities. However, the greedy performs best throughout most of the horizon. If additional resources are received throughout in the planning horizon, the greedy policy might outperform the others. If the reemerging outbreak is more explosive, the greedy policy runs out of resources and quickly increase the number of fatalities. Although the naive policy outperforms the ADP policy in some scenarios, the ADP policy better adjusts to the actual realizations of the dispersal rate, and thus, on average, outperforms the naive policy.

9.5 Sensitivity Analysis

The following sections varies key parameters for the resource allocation model. Section 9.5.1 investigates the effect of reduced kit sizes for the medical resources. The analysis

is performed on the base case on the 2010 Haiti outbreak data, due to the kit size not being connected to the epidemic evolution. In Section 9.5.2, the arrival and availability of medical resources are varied and investigated. Because the greedy approach seemed to contain the outbreak when using the epidemic model based on the 2010 Haiti outbreak, the more unstable, alternative epidemic model is employed. For the same reason, the alternative epidemic model is used when various bacteria dispersal probability distributions are investigated in Section 9.5.3. Lastly, Section 9.5.4 investigates the performance when the planning horizon is extended.

9.5.1 Reduction in resource kit size

In the base case, the resources are sent in kits of 1000. As explained in Section 7.3, the demand are ceiled, e.g. if the demand for ORS in a region is 1600, two kits will be sent. Treatment will be available for 2000 throughout the week, thus if the demand increases up to 2000, it can still be satisfied. However, this also implies that too many resources may be sent. The smaller the kit size, the more accurate the demand is met, but it may also increase the computational time of training the VFA and employing the ADP policy. Too small kit sizes are unrealistic, because sending single doses of for example vaccines, would be too expensive and impractical. In the following sections reductions in kit sizes are investigated. Although an increase in kit size is also possible, it would be unreasonable in practice. With an increase to, for instance 10 000 resources per kit, allocating and utilizing a single resource would exceed the capacity in some regions.

The mean performance of the policies with kit sizes of 500 and 100 are shown in Figure 9.12a and Figure 9.12b, respectively. There is a noticeable difference in performance compared to a kit size of 1000, seen in Figure 9.5. With a decreasing kit size, the naive policy performs better, while both the ADP policy and the greedy policy performs worse. The solution time, not including the time it takes to train the VFA, for the different kit sizes are reported in Table 9.5. As expected, a decreasing kit size will increase the solution time. However, the increase in solution time may be neutralized by employing more sophisticated step sizes in the local search.

With a kit size of 500, the fatalities with the greedy and naive approach are similar, as seen in Figure 9.12a. Given a realization path of the expected dispersal rate of 0.025 every week, the cholera-induced fatalities are about 15 000 and about 16 000 for the greedy and naive policy, respectively. The decision employing the greedy and naive policies are shown in Table 9.6. Although the total number of fatalities employing the different policies are similar, the allocation differs significantly. While the greedy algorithm allocates the vaccines immediately, the naive approach is still distributing them at the beginning of week 10. Even though the total fatalities are similar, the allocations between the policies differ with a kit size of 500.

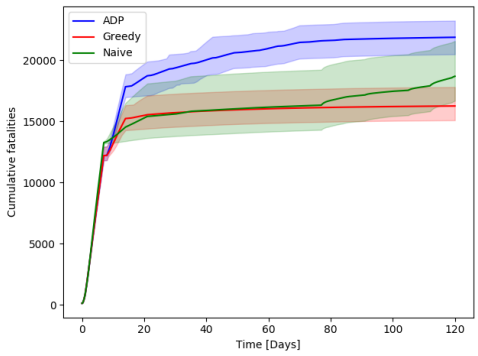

(a) Kit size 500

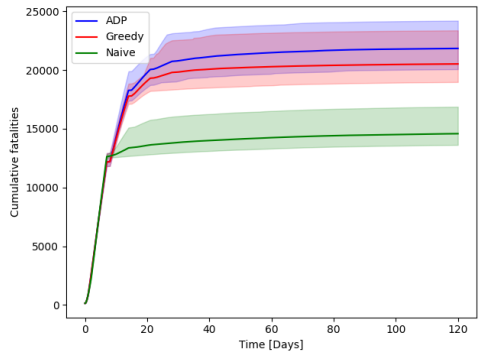

(b) Kit size 100

Figure 9.12: Mean performance in terms of cumulative cholera-induced fatalities and 95th and 5th percentiles after 100 simulations, for various policies with kit size (a) 500 and (b) 100.

Table 9.5: Best, mean and worst solution time in seconds for the ADP policy on 100 different dispersal realizations, with various kit sizes.

Kit size	Metric	Solution time
1000 (base case)	Best	53s
	Mean	83s
	Worst	106s
500	Best	87s
	Mean	102s
	Worst	151s
100	Best	88s
	Mean	184s
	Worst	217s

Table 9.6: Resource allocations for different policies (Greedy / Naive), with kit size of 500 and expected dispersal rate as realization path for selected weeks. The resources are in multiples of 500.

Region	Week 2				Week 10			
	Vac.	Dis.	ORS	Ant.	Vac.	Dis.	ORS	Ant.
Ouest	171 / 38	0 / 34	3 / 19	3 / 9	0 / 7	0 / 6	1 / 3	1 / 1
Sud-Est	57 / 8	0 / 7	1 / 4	1 / 2	0 / 5	0 / 4	1 / 3	1 / 2
Nord	57 / 9	0 / 8	1 / 5	1 / 3	0 / 5	0 / 4	1 / 2	1 / 1
Nord-Est	57 / 1	0 / 1	1 / 0	1 / 0	0 / 2	0 / 2	1 / 1	1 / 0
Artibonite	229 / 40	0 / 35	4 / 20	4 / 10	0 / 5	0 / 4	1 / 2	1 / 1
Centre	114 / 11	0 / 10	2 / 5	2 / 2	0 / 1	0 / 2	1 / 1	1 / 1
Sud	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0
Grande Anse	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0
Nord-Ouest	57 / 8	0 / 7	1 / 4	1 / 2	0 / 3	0 / 2	1 / 1	1 / 0
Nippes	58 / 0	0 / 0	1 / 0	1 / 0	0 / 0	0 / 0	1 / 0	1 / 0
Total	800 / 115	0 / 102	14 / 57	14 / 28	0 / 28	0 / 24	8 / 13	8 / 6

Interestingly, in Figure 9.12b, both the ADP and the greedy policy perform worse, but the naive policy performs better, compared to the original kit size of 1000, shown in Figure 9.5. Even though the total fatalities when employing the naive and greedy policies are similar, the allocations are dissimilar, the improvement of the naive policy and the decline of the ADP and greedy policy can be explained as independent phenomena. A possible explanation for the greedy and ADP policies performing worse is the buffering effect of ceiling the demand. When the kit size is large, the demand is not only covered, but an larger buffer is included for the following week. As the kit size is reduced, the buffer is reduced, and the performance worsens. This cannot, however, explain the improvement of the naive policy. The naive policy allocates resources based on an infection ratio across the planning horizon. Ideally, each region would receive the same resource ratio as its infection ratio. However, this is not always feasible, because the resources are integer. When the kit size is large, the number of integer resources are low, but as the kit size decreases, it is easier to allocate resources closer to the infection ratio. For instance, given 10 available vaccines, the number of vaccine resource kits would be 10 with a kit size of 1 and 1 with a kit size of 10. In this case it is easier to divide 10 than 1 fairly among three recipients. In the former case all three recipients receive at least three resources. In the latter, only one recipient receive a kit, leaving nothing left for the others.

9.5.2 Impact of availability of resources and arrival time

When an epidemic outbreak has occurred, a fast response is essential. Even in a region where outbreaks occur seasonally, the response resources might not be immediately available. Vaccines, being allocated from a global stockpile, must be applied for, as described in Section 2.4.4. In the base case, the vaccines arrive and are allocated two weeks after the outbreak began. This section investigates the effect the arrival time of vaccines and other intervention resources have on the alternative epidemic outbreak. The effect of various availability of resources are also investigated, such as an increase in available vaccines or rehydration solutions.

Vaccination: The effects of receiving vaccines earlier than in the base case are shown in Table 9.7. The effect of early vaccination seems to be negligible for the ADP, naive and myopic policies. However, it seems to be significant for the greedy policy. Investigating the individual decisions, the allocation of vaccines is the same for the greedy and the ADP policy. The vaccine allocations are the same, however, the effect of them are different, suggesting that the allocations of other intervention resources impact the benefit of earlier vaccination. The base case analysis revealed that the ADP and the naive policies allocate more resources later in the planning horizon, to better adjust if the epidemic reemerges. The greedy policy allocates to satisfy the immediate demand as much as possible, and is thus more sensitive to reemergence of the epidemic. The impact of a reemergence is reduced the earlier the vaccines are allocated, thus the marginal benefit of early vac-

cination allocation is higher for the greedy policy, compared to the ADP and the naive policies. This is a possible explanation for the noticeable improvement for greedy policy with early availability of vaccines, compared to the ADP and naive policies, particularly in the worst-case scenarios.

Table 9.7: Best, mean and worst performance of policies for various scenarios of resource availability and arrival for 100 epidemic simulations. Number of fatalities in thousands.

Resource scenario	Metric	ADP	Greedy	Naive	Myopic
Two week vaccines (base case)	Best	74	85	89	171
	Mean	100	132	107	178
	Worst	160	176	119	185
One week vaccines	Best	70	86	87	170
	Mean	98	125	106	178
	Worst	159	167	128	183
Immediate vaccines	Best	69	78	89	170
	Mean	100	115	107	177
	Worst	161	158	118	179

Although the arrival time of the vaccines, did not have a significant impact on the ADP and the naive policies, the increase in vaccines did. The performance for various number of vaccines are shown in Table 9.8. The overall trend is a reduction in fatalities when the availability of vaccines increase.

In the 600 000 vaccine scenario, the best performance for the ADP does not change. However, the performance for the greedy policy change noticeably. As 800 000 vaccines are available, the performance is improved further for all policies. Lastly, when 1 million vaccines are available, all performances improve, except the greedy worst-case. However, that could be due to an outlier realization path not explored during the other vaccine cases. The consistent improvement in performance, regardless of the additional amount of vaccines, indicates that a high availability of vaccines, even if they arrive two weeks after the initial outbreak, may provide significant improvement in containment. In terms of vaccination strategy employed, the ADP, the greedy and the naive policies all employ a reactive strategy, allocating vaccines to regions already having outbreaks. The greedy and naive do this by default, but the ADP does not change the strategy through its local search.

Overall, the number of available vaccines seems to reduce fatalities more than earlier arrival of vaccines. When applying to the International Coordinating Group on Vaccine Provision (ICG) for vaccines, the focus should thus be on collecting reliable surveillance data on infections and spread to ensure as many vaccines as possible, instead of applying as quickly as possible to get the vaccines earlier.

Rehydration solution: In Section 9.4 it was identified that a possible explanation for the ADP and naive policy performing better than the greedy is that the greedy policy

Table 9.8: Best, mean and worst performance of policies for various scenarios of resource availability and arrival for 100 epidemic simulations. Number of fatalities and vaccines in thousands.

Resource scenario	Metric	ADP	Greedy	Naive	Myopic
400 vaccines (base case)	Best	74	85	89	171
	Mean	100	132	107	178
	Worst	160	176	119	185
600 vaccines	Best	80	78	83	164
	Mean	93	116	101	171
	Worst	133	164	112	180
800 vaccines	Best	66	68	76	157
	Mean	86	99	94	165
	Worst	129	134	105	176
1000 vaccines	Best	63	66	68	150
	Mean	80	91	88	160
	Worst	115	145	100	172

depletes its rehydration solution resources the earliest. ORS is essential to treat symptomatic cholera-infections, thus the depletion cause a rise in fatalities. If that is the case, increasing the number of ORS treatment should significantly decrease the fatalities when employing the greedy policy. This does seem to be the case for the greedy policy, when increasing the available ORS with 100 000 units to 300 000 units. For the mean- and worst-case performance, the fatalities are reduced from 132 000 to 112 000 and from 176 000 to 143 000, respectively, as seen in Table 9.9. The best-case performance only improves marginally for the greedy policy. A possible explanation is that the reemerging outbreak is sufficiently small to be covered with the original amount of ORS. The increase in fatalities is then explained by an increase in infections, that even with treatment, will increase the number of fatalities somewhat. Increasing the ORS with 200 000 units to a total of 400 000 units does not seem to reduce fatalities further, when employing the greedy policy.

Table 9.9: Best, mean and worst performance of policies for various scenarios of resource availability and arrival for 100 epidemic simulations. Number of fatalities and rehydration solutions in thousands.

Resource scenario	Metric	ADP	Greedy	Naive	Myopic
200 ORS (base case)	Best	74	85	89	171
	Mean	100	132	107	178
	Worst	160	176	119	185
300 ORS	Best	81	77	119	171
	Mean	102	112	136	179
	Worst	137	143	158	184
400 ORS	Best	84	81	133	171
	Mean	110	114	158	178
	Worst	147	145	188	184

Similar to the greedy policy, the ADP policy improves its worst-case performance when an additional 100 000 ORS are available. However, when it is further increased to 400

000 ORS treatments, the performance is actually worse. This could plausibly be explained by statistical inaccuracy from 100 simulations. Table 9.10 shows the aggregated allocation of ORS in the worst-case performance with 200 000 and 400 000 available ORS. Note that the realization paths of the dispersal rate are not necessarily the same, between resource cases. In the beginning the ORS allocations are similar, but as the epidemic progresses, and particularly in the last weeks of the planning horizon, the ADP allocates more ORS in the additional ORS scenario. When the epidemic reemerges, which it does in both the worst-case performances, the ADP policy does not have to ration the resources in the additional ORS scenario, but can instead allocate in accordance with demand, utilizing 251 ORS kits instead of 200. Thus, the decrease in fatalities for the worst-case performance is in fact due to the increase in ORS resources.

Table 9.10: Worst-case aggregated ORS allocation for the ADP policy for the base case with 200 000 ORS treatments and the case with additional ORS, i.e. 400 000 ORS treatments.

Week	Base case	Additional ORS
0	2	2
1	9	9
2	12	12
3	13	13
4	14	14
5	13	12
6	12	13
7	10	13
8	12	10
9	9	12
10	11	12
11	15	16
12	15	19
13	20	23
14	17	22
15	15	10
16	1	17
17	0	22
Total	200	251

An unexpected result is the performance for the naive policy, which actually performs worse when additional ORS is available. The difference is too large to be caused by statistical inaccuracy. When investigating the particular allocations made, the naive policy fails to allocate any ORS to Sud, Grande Anse and Nippes early, because those regions have late or no outbreaks when using the calibrated bacteria dispersal rate. The ADP policy does allocate to these regions, proving they have more critical outbreaks when the dispersal rate is stochastic. This explains why the ADP policy ensures less fatalities than the naive policy, but not why the naive policy performs worse. The regions having outbreaks with a deterministic, calibrated dispersal rate receive more ORS. The increase could cause certain infected individuals to survive longer, excreting more bacteria, which in turn cause more explosive outbreaks in the regions that are never responded to using

the naive policy. While saving lives in the regions with more fatalities, the increased bacteria concentration in regions that do not receive any ORS cause a net increase in fatalities.

An increase in the ORS can improve the performance, particularly for the greedy policy. The ADP is not as sensitive to reemerging epidemics, due to rationing the ORS earlier in the planning horizon, but in the most explosive reemerging epidemics, additional ORS also improves the ADP policy performance. For both the ADP and greedy policies, an additional 100 000 ORS treatments seems to be sufficient.

9.5.3 Dispersal probability distribution

Cholera outbreaks have been tied to environmental fluctuations (Olson et al., 2018, pp. 12). During drier periods, the dispersal of bacteria through river networks could be close to zero, causing the original outbreak region to be even more severely affected, but reducing the spread to neighboring regions. During periods of particularly large rainfalls, the bacteria dispersal by river networks might be drastically higher than normal, dampening the outbreak in the originating region, but increasing the probability of outbreaks in other regions. As seasonal outbreaks occur, the magnitude of the outbreaks can be tied to weather data and used to obtain better estimates of the dispersal rate probability distribution. This section investigates the effect of various bacteria dispersal distributions. The different distributions used are shown in Table 9.11. Note that the calibrated value distribution is deterministic and based on the dispersal rate reported in Bertuzzo et al. (2011). Projected cholera-induced fatalities under various policies for the different dispersal distributions are summarized in Table 9.12.

Table 9.11: Possible weekly dispersal rates and their probability of occurring for the different dispersal distributions.

Dispersal distribution	Metric	Value	Expected value
Base case distribution	Dispersal rate Probability	(0.0, 0.025, 0.25) (0.25, 0.50, 0.25)	0.075
Calibrated value	Dispersal rate Probability	(0.025) (1.0)	0.025
Low variance	Dispersal rate Probability	(0.0, 0.025, 0.25) (0.10, 0.80, 0.10)	0.045
High variance	Dispersal rate Probability	(0.0, 0.025, 0.25) (0.333, 0.333, 0.333)	0.092
Low maximum dispersal	Dispersal rate Probability	(0.0, 0.025, 0.050) (0.25, 0.50, 0.25)	0.025

The distributions with the lowest expected values, the *calibrated value* and the *low maximum dispersal distribution*, seem to have the lowest fatalities, indicating that a lower dispersal rate is preferable. A low dispersal rate, for instance due to drought, may result in more explosive outbreaks, but isolated to the regions they originated in. The results

Table 9.12: Mean, best and worst performance of policies for various scenarios of bacteria dispersal rate distributions across 100 simulations. The calibrated value scenario is deterministic. Number of fatalities in thousands.

Dispersal distribution	Metric	ADP	Greedy	Naive	Myopic
Base case distribution	Best	74	85	89	171
	Mean	100	132	107	178
	Worst	160	176	119	185
Calibrated value	-	81	88	94	180
Low variance	Best	71	82	89	172
	Mean	88	112	101	179
	Worst	153	171	116	184
High variance	Best	77	87	89	171
	Mean	110	135	110	177
	Worst	166	170	120	185
Low maximum dispersal	Best	68	76	89	179
	Mean	87	89	98	181
	Worst	108	151	121	184

in Table 9.12 indicates that it is more efficient to respond to explosive outbreaks in a single or a few regions, compared to smaller outbreaks in many regions. As the expected dispersal rate increases, the fatalities also consistently increase. This is likely due to the instability resulting from a high bacteria excretion rate. If bacteria are dispersed to several regions early, a high bacteria concentration is quickly built up, ensuring more sources for future spread, but also enough excretion to sustain a high concentration and create outbreaks in the regions already having symptomatic infections.

The ADP policy ensures on average less fatalities compared to all other policies investigated across all distributions. The best-case performance is also consistently better for the ADP policy. If working correctly, the VFA in the ADP policy should be able to account for future possible dispersal that could affect the performance, and aims to minimize the expected number of fatalities. Since the average performance is consistently best for the ADP policy across various distributions, it seems that the VFA does capture and accounts for the stochasticity in the dispersal rate.

The ADP and greedy policies seem to be more dependent on the dispersal distribution than the naive and myopic policies. A possible explanation for this is that the ADP and greedy policy are better at containing the outbreaks early. Thus, if the dispersal rate is low, the likelihood of reemergence is low, and large allocations early on are the most effective strategies. The ADP approach aims to find the expected cost, thus balancing the likelihood of reemergence with the effect of containment, explaining its dependency on dispersal distribution.

The mean performance for the greedy policy seem more sensitive to the distributions than the ADP policy. For instance, when going from the deterministic calibrated value case to the high variance case, the fatalities increase by 36%, from 81 000 to 110 000, for the

ADP policy and 53%, from 88 000 to 135 000, for the greedy policy. The ADP policy is trained on the distribution, thus learning the probability of larger reemerging outbreaks occurring and accounting for that when allocating resources early in the planning horizon. Therefore, the higher adaptability for the ADP policy compared to the greedy policy is expected. The ADP policy performance is also relatively consistent in the worst-case scenario. However, when the maximum dispersal rate is low, the performance is significantly better than with the other distributions in the worst-case scenario. This indicates that the ADP policy performs better when responding to outbreaks in fewer regions or when the outbreaks in other regions are more gradual.

The expected value of the dispersal rate is higher than the calibrated value for the base case distribution, low variance distribution and the high variance distribution, as reported in Table 9.11. For the base case and high variance distribution, the naive policy performs better than the greedy. The fatalities employing the greedy policy is reduced when the expected dispersal rate decreases, as is the case in the low variance distribution, the calibrated value distribution and the low maximum distribution. In Section 9.4, an explanation for the great decline in performance of the greedy policy late in the planning horizon was due to a reemerging outbreak. If the dispersal rate is higher, more bacteria can be transported in shorter time, thus the outbreaks can reemerge more rapidly and be more challenging to contain. However, the greedy policy consistently has a better performance in the best-case, compared to the naive policy, and almost a consistently worse performance in the worst-case, except for when the maximum dispersal is low or always equal to the calibrated value. A low dispersal realization path is possible in every distribution, while a high is possible in every distribution except the calibrated and low maximum distribution. Thus, the performances indicate that the naive policy is more robust to higher dispersal rates compared to the greedy policy.

When the dispersal rate is low, the outbreaks in specific regions can be more explosive, and is one explanation why drought can cause cholera outbreaks. However, the ADP and the greedy policies seem to perform better when the bacteria dispersal between regions is low. The ADP policy shows some adaptability, as expected. However, the performance of the remaining policies seems dependent on the dispersal distributions, thus getting better estimates of the true distribution is worthwhile. For instance, if choosing between the greedy and the naive policies, the naive is more robust to reemerging outbreaks, and thus preferable when the expected dispersal rate is high. However, if the probability of reemerging outbreaks are low, the greedy policy is superior to the naive.

9.5.4 Increased planning horizon

A cholera outbreak typically lasts between 2 and 4 months, as described in Section 2.3. The alternative cholera model causes the outbreak to reemerge. Given the four-month

planning horizon employed in the base case, the ADP algorithm seems to perform the best, as seen in Figure 9.9. However, if the epidemic reemerges, the evaluation of resource allocation policies becomes more difficult. Since the epidemic reemerges, one could apply the model again, with a new initial response time $t = 0$, arguing that the two outbreaks are distinct. However, the response to the first outbreak might impact the magnitude of the reemerging outbreak, therefore, evaluating the policy choice jointly by increasing the planning horizon may be more reasonable.

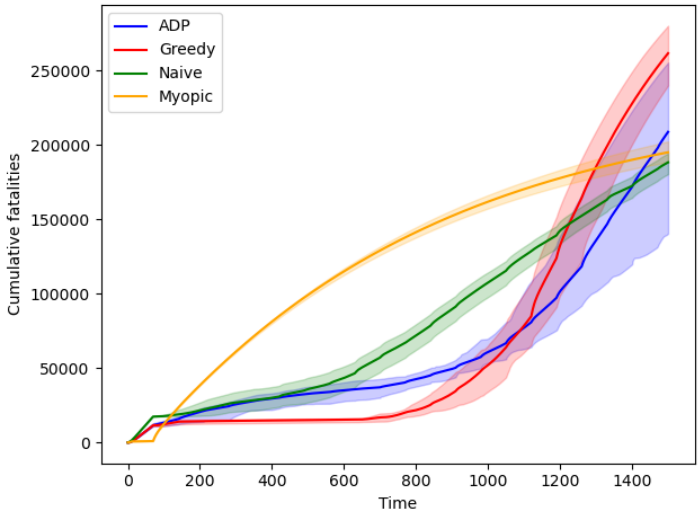

Figure 9.13: Comparison of policies with mean cumulative fatalities across 100 epidemic realizations when planning horizon is 150 days. Intervals are 95th and 5th percentiles of the respective policy.

The evaluation for the various policies employed on the alternative epidemic outbreak with the base case resources for a planning horizon of 150 days, or about five months, is shown in Figure 9.13. Observe that the ADP and greedy policies have increased fatalities as the outbreak reemerges. As previously identified, due to saving more resources early on, the ADP policy better responds to the reemergence. However, if the reemerging outbreak is large, the naive policy eventually surpasses the performance of the ADP policy. If the reemerging outbreak is smaller, the ADP policy saved a sufficient amount of resources to keep the fatalities relatively low and performs significantly better than all other policies, as seen in the lower interval of the ADP policy in Figure 9.13.

A notable case is the myopic policy, which, after depleting its resources early, has a drastic increase in fatalities, which eventually stagnates. This can be explained by the increased excretion rate. By providing oral rehydration solutions, the infected individuals keep introducing more bacteria to the water reservoir. However, if there is a lot

of fatalities early on, less bacteria is introduced, and the bacteria decay eventually contains the epidemic. In the case with higher excretion, the model seems to fail to account for a boundary condition ending the epidemic. Additional disinfectants and antibiotics could mitigate the effect, by reducing the bacteria concentration in the water reservoir and decrease the excretion rate, respectively.

Although the excretion rate is within reasonable bounds determined from epidemiological literature, the set of parameters selected may together not be sufficiently realistic. For instance, the relative cholera-induced fatality rate when not receiving rehydration treatment χ, estimated to be 250, could be too high. If the parameter is decreased, symptomatic individuals not receiving treatment would survive for longer, increasing their bacteria contribution to the water reservoir. This would in turn increase the infection rate and increase the likelihood for reemerging outbreaks even if the fatalities are high earlier, thus negatively impacting the myopic performance. Another possible explanation is the upper bound of the proportion of the population ingesting contaminated water. If no disinfectants are allocated, this is assumed to be the entire population. However, some individuals may have sanitation infrastructure with advanced filtering systems, ensuring cholera-free water even without disinfectants. By reducing the upper bound of the proportion ingesting contaminated water, the reemerging outbreak would not be as explosive, and possibly stagnate and stabilize earlier, i.e. behaving more as expected.

Chapter 10

Concluding Remarks

This chapter concludes this thesis and identifies opportunities for improvement and future research. Section 10.1 presents the thesis conclusion, and the future research opportunities are identified and discussed in Section 10.2.

10.1 Conclusion

The purpose of this thesis was to develop decision-support tools for responses during epidemic outbreaks for diseases with known treatment methods in financially weak regions and to investigate the efficiency of various policies during such outbreaks. To provide this decision-support, a combined epidemic and resource allocation model for multiple regions and multiple intervention methods was developed.

The decision-support is aimed at the allocation of medical resources, while also satisfying the medical personnel and temporary medical facility constraints. The literature review revealed that there is a lack of decision-support models combining multi-region and multi-intervention epidemic and resource allocation models, as well as a lack of cholera-modeling within an approximate dynamic programming (ADP) framework. The proposed epidemic model is a Susceptible-Asymptomatic-Infected-Recovered-Bacteria (SAIR-B) model for cholera. The model combines previous works on spatially explicit cholera models with intervention modeling for epidemic diseases, and extends the work by including oral rehydration solutions (ORS) as an additional intervention parameter. Furthermore, the proposed resource allocation model is developed within an ADP framework to account for the effect of immediate decisions on future epidemic spread, thus providing a holistic view, balancing the need to contain the epidemic early with possible

reemerging outbreaks later. The resources allocated are: vaccines, disinfectants, rehydration solutions and antibiotics.

The ADP resource allocation model employs a neural network as its value approximation function. Using realistic, simulated epidemic outbreaks, the neural network aims to learn the value, or cost, of making a decision, for future disease-induced fatalities throughout the epidemic. The model can be used to simulate various policies before an epidemic, or it can be used as a decision-support tool during an epidemic, updating the parameters in both the cholera model and the epidemic model as the outbreak progresses. In addition to the ADP approach, three other policies were developed for comparison: a greedy, a myopic and a naive one.

In the computational study, the performance of the different epidemic response policies in various situations is investigated. The results indicate that the ADP and greedy policies are sensitive to the bacteria dispersal distributions, but if the excretion rate is high, the ADP policy is robust and consistently ensures the least fatalities. However, when the bacteria excretion rate is low, the probability of reemergence is significantly lower, and the greedy policy performs best. In that case, the ADP policy also performs well, although not as good as the greedy. Overall, considering the uncertainty in the dispersal rate, the ADP policy seems the best policy option among the investigated policies, when employing the alternative epidemic model.

Resources are sent in kits, because sending individual resources, e.g. vaccines, would be both impractical and expensive. Although one might expect that smaller kit sizes would improve performance due to meeting the demand more accurately, the fatalities when employing an ADP or greedy policy increase with smaller kit sizes, indicating a buffer effect. A smaller kit size improved the naive allocation policy, likely due to increased accuracy in meeting the demand. By adjusting the initial condition for regions with a high likelihood of infections, even if there are no specific suspected cases there yet, may mitigate the need for a buffer.

Additional ORS did not improve the average and best-case performance for the ADP and greedy policies. However, if the epidemic reemerged, and the second outbreak was large, both the ADP and greedy policies would eventually deplete their ORS resources. In such cases, some additional ORS resources would improve performance. Therefore, having a buffer stock of ORS and ensuring a functioning replenishment of it throughout the outbreak is important. Nevertheless, in order to further decrease the fatalities, vaccines should be the main focus.

The International Coordinating Group for Vaccine Provision (ICG) manages a global stockpile of oral cholera vaccines, distributing them as epidemic outbreaks occur. Because applying for vaccines occurs after the outbreak has begun, the decision-makers might be tempted to apply as early as possible to ensure the earliest possible arrival of

the vaccines. However, additional vaccines seemed to have a more significant impact on reducing the total number of cholera-induced fatalities compared to earlier arrival. Therefore, it is advised to focus on sufficient surveillance of the outbreak to convince the ICG of a larger allocation of vaccines, instead of applying as early as possible.

10.2 Future Research

The objective of the resource allocation model is to minimize the total number of cholera-induced fatalities throughout the planning horizon. The objective is evaluated by numerically solving the epidemic model, giving rise to a mixed-integer nonlinear programming (MINLP) problem. Although the ADP policy appeared to converge towards a policy, it solved the resource allocation problem at each time period heuristically. In the Haiti base case, the greedy approach consistently performed the best. If the proposed ADP policy worked ideally, and the greedy policy is in fact optimal in the Haiti base case, the ADP policy should have converged towards the greedy policy, instead of saving resources for later. Thus, employing more sophisticated heuristics than the local search procedure to solve the resource allocation problem for each time period, may further improve the performance of the ADP approach.

Comprehensive background research was conducted to arrive at the formulated operation constraints presented in Section 6.2. The case study performed in this thesis focused on the 2010 Haiti outbreak. Ignoring distribution time to the regions in Haiti may be a reasonable assumption in a small country or region, but as the size of the country increases, distribution time should eventually be accounted for, and can be included in a future formulation. Furthermore, the medical personnel is assumed to be homogeneous, e.g. there is no differentiation between physicians and nurses. This formulation assumes that it is more important that casualties receive help, than who administers it. However, differentiating between medical professions may be more realistic and thus included in a future formulation.

Instead of relying too heavily on parameters from the literature, a more extensive calibration with higher degrees of freedom may be conducted on the 2010 Haiti outbreak or other epidemic outbreaks. Such a calibration may provide more realistic simulations for the resource allocation to learn from, increasing the validity of the model. However, the purpose should be to provide realistic parameter values to learn more about response to future outbreaks, not to perfectly fit historical outbreaks. If data on disease-induced fatalities are available, the relative increase in fatality rate when not receiving rehydration solutions may also be better estimated. Additionally, when increasing the planning horizon, there were realization paths using the alternative epidemic model with increased excretion rate where the reemerging outbreaks were significant. A possible explanation and future improvement of the model is the assumption regarding disinfectants. The assumed

proportion ingesting contaminated water to some degree, when no disinfectants are allocated, is the entire population. Even with sanitation infrastructure, water sources can get contaminated and cause cholera, but by including the proportion ingesting contaminated water in the epidemic model calibration, an upper bound may be found, resulting in more realistic epidemic simulations.

Cholera outbreaks can be caused by environmental fluctuations and findings in the sensitivity analysis suggested that different bacteria dispersal distributions greatly affected the different policies' performance, measured as the number of cholera-induced fatalities. The dispersal distribution is a key parameter to the resource allocation model, and increasing the certainty of which the bacteria disperse under various climatic conditions is worth considering. For instance, the dispersal rate may be correlated with local rainfall data in regions with seasonal cholera outbreaks.

Finally, the generalizability of the findings to other diseases than cholera should be investigated. Although the modeling framework presented in this thesis could be extended to other communicable diseases, it would require different epidemic models, and possibly other intervention resources, depending on the disease of interest. Therefore, to test the generalizability of the results of this thesis, the findings on the vaccine provision from the ICG could be investigated for diseases like meningitis and yellow fever, for which the ICG also have vaccine stockpiles.

Bibliography

Allen, E., 2016. Environmental variability and mean-reverting processes. Discrete & Continuous Dynamical Systems-B 21, 2073.

Allen, L.J., 2017. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. Infectious Disease Modelling 2, 128–142.

Altay, N., Green III, W.G., 2006. Or/ms research in disaster operations management. European journal of operational research 175, 475–493.

Andrews, J.R., Basu, S., 2011. Transmission dynamics and control of cholera in haiti: an epidemic model. The Lancet 377, 1248–1255.

Anparasan, A., Lejeune, M., 2017. Resource deployment and donation allocation for epidemic outbreaks. Annals of Operations Research , 1–24.

Arora, H., Raghu, T., Vinze, A., 2010. Resource allocation for demand surge mitigation during disaster response. Decision Support Systems 50, 304–315.

Azaele, S., Maritan, A., Bertuzzo, E., Rodriguez-Iturbe, I., Rinaldo, A., 2010. Stochastic dynamics of cholera epidemics. Physical Review E 81, 051901.

Becker, N.G., Starczak, D.N., 1997. Optimal vaccination strategies for a community of households. Mathematical Biosciences 139, 117–132.

Bertuzzo, E., Azaele, S., Maritan, A., Gatto, M., Rodriguez-Iturbe, I., Rinaldo, A., 2008. On the space-time evolution of a cholera epidemic. Water Resources Research 44.

Bertuzzo, E., Mari, L., Righetto, L., Gatto, M., Casagrandi, R., Blokesch, M., Rodriguez-Iturbe, I., Rinaldo, A., 2011. Prediction of the spatial evolution and effects of control measures for the unfolding haiti cholera outbreak. Geophysical Research Letters 38.

Boonmee, C., Arimura, M., Asada, T., 2017. Facility location optimization model for emergency humanitarian logistics. International Journal of Disaster Risk Reduction 24, 485–498.

Brandeau, M.L., Zaric, G.S., Richter, A., 2003. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. Journal of health economics 22, 575–598.

Büyüktahtakın, İ.E., des Bordes, E., Kıbış, E.Y., 2018. A new epidemics–logistics model: Insights into controlling the ebola virus disease in west africa. European Journal of Operational Research 265, 1046–1063.

Capasso, V., Paveri-Fontana, S., 1979. A mathematical model for the 1973 cholera epidemic in the european mediterranean region. Revue d'épidémiologie et de Santé Publiqué 27, 121–132.

Caunhye, A.M., Nie, X., Pokharel, S., 2012. Optimization models in emergency logistics: A literature review. Socio-economic planning sciences 46, 4–13.

Centers for Disease Control and Prevention, 2019. Treatment. URL: `https://www.cdc.gov/vhf/ebola/treatment/index.html`.

Checchi, F., 2009. Principles of infectious disease transmission.

Chin, C.S., Sorenson, J., Harris, J.B., Robins, W.P., Charles, R.C., Jean-Charles, R.R., Bullard, J., Webster, D.R., Kasarskis, A., Peluso, P., et al., 2011. The origin of the haitian cholera outbreak strain. New England Journal of Medicine 364, 33–42.

Codeço, C.T., 2001. Endemic and epidemic dynamics of cholera: the role of the aquatic reservoir. BMC Infectious diseases 1, 1.

Colwell, R.R., 1996. Global climate and infectious disease: the cholera paradigm. Science 274, 2025–2031.

Coppola, D.P., 2006. Introduction to international disaster management. Elsevier.

Coşgun, Ö., Büyüktahtakın, İ.E., 2018. Stochastic dynamic resource allocation for hiv prevention and treatment: An approximate dynamic programming approach. Computers & Industrial Engineering 118, 423–439.

Crooks, A.T., Hailegiorgis, A.B., 2014. An agent-based modeling approach applied to the spread of cholera. Environmental Modelling & Software 62, 164–177.

Dandekar, R., Barbastathis, G., 2020. Quantifying the effect of quarantine control in covid-19 infectious spread using machine learning. medRxiv .

Dasaklis, T.K., Pappis, C.P., Rachaniotis, N.P., 2012. Epidemics control and logistics operations: A review. International Journal of Production Economics 139, 393–410.

Eisenberg, M.C., Kujbida, G., Tuite, A.R., Fisman, D.N., Tien, J.H., 2013. Examining rainfall and cholera dynamics in haiti using statistical and dynamic modeling approaches. Epidemics 5, 197–207.

Emch, M., Feldacker, C., Islam, M.S., Ali, M., 2008. Seasonality of cholera from 1974 to 2005: a review of global patterns. International journal of health geographics 7, 31.

Gazi, N., Das, K., Mukandavire, Z., Chiyaka, C., Das, P., 2010. A study of cholera model with environmental fluctuations. International Journal of Mathematical Models and Methods in Applied Sciences 4, 150–155.

Global Task Force on Cholera Control, 2004. Cholera outbreak: assessing the outbreak response and improving preparedness.

Global Task Force on Cholera Control, 2017. Ending cholera. a global roadmap to 2030.

GNS Science, 2020. How long does an earthquake last? URL: https://www.gns.cri.nz/Home/Learning/Science-Topics/Earthquakes/Monitoring-Earthquakes/Other-earthquake-questions/How-long-does-an-earthquake-last.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep learning. MIT press.

Gupta, S., Starr, M.K., Farahani, R.Z., Matinrad, N., 2016. Disaster management from a pom perspective: Mapping a new domain. Production and Operations Management 25, 1611–1637.

Hartley, D.M., Morris Jr, J.G., Smith, D.L., 2005. Hyperinfectivity: a critical element in the ability of v. cholerae to cause epidemics? PLoS Med 3, e7.

Havumaki, J., Meza, R., Phares, C.R., Date, K., Eisenberg, M.C., 2019. Comparing alternative cholera vaccination strategies in maela refugee camp: using a transmission model in public health practice. BMC Infectious Diseases 19, 1–17.

Institut Haitien e Statistique et d'Informatique, 2015. Population totale, de 18 ans et plus. URL: https://web.archive.org/web/20151106110552/http://www.ihsi.ht/pdf/projection/Estimat_PopTotal_18ans_Menag2015.pdf.

Kanchanaraksa, S., 2008. Lecture notes on epidemiologic investigation.

Kaper, J.B., Morris, J.G., Levine, M.M., 1995. Cholera. Clinical microbiology reviews 8, 48–86.

Kermack, W.O., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character 115, 700–721.

King, A.A., Ionides, E.L., Pascual, M., Bouma, M.J., 2008. Inapparent infections and cholera dynamics. Nature 454, 877–880.

Kovacs, G., Moshtari, M., 2019. A roadmap for higher research quality in humanitarian operations: A methodological perspective. European Journal of Operational Research 276, 395–408.

Lemos-Paião, A.P., Silva, C.J., Torres, D.F., 2017. An epidemic model for cholera with optimal control treatment. Journal of Computational and Applied Mathematics 318, 168–180.

Leshno, M., Lin, V.Y., Pinkus, A., Schocken, S., 1993. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. Neural networks 6, 861–867.

Liu, M., Zhang, Z., Zhang, D., 2015. A dynamic allocation model for medical resources in the control of influenza diffusion. Journal of Systems Science and Systems Engineering 24, 276–292.

Liu, Q., Jiang, D., Hayat, T., Alsaedi, A., 2019. Dynamical behavior of a stochastic epidemic model for cholera. Journal of the Franklin Institute .

Long, E.F., Nohdurft, E., Spinler, S., 2018. Spatial resource allocation for emerging epidemics: A comparison of greedy, myopic, and dynamic policies. Manufacturing & Service Operations Management 20, 181–198.

Ludkovski, M., Niemi, J., 2010. Optimal dynamic policies for influenza management. Statistical Communications in Infectious Diseases 2.

Mari, L., Bertuzzo, E., Righetto, L., Casagrandi, R., Gatto, M., Rodriguez-Iturbe, I., Rinaldo, A., 2012. Modelling cholera epidemics: the role of waterways, human mobility and sanitation. Journal of the Royal Society Interface 9, 376–388.

Mete, H.O., Zabinsky, Z.B., 2010. Stochastic optimization of medical supply location and distribution in disaster management. International Journal of Production Economics 126, 76–84.

MSPP, 2020. Centre de documentation. URL: `https://mspp.gouv.ht/newsite/documentation.php`.

MSPP and CDC, 2011. Haiti cholera training manual: A full course for healthcare providers.

Mwasa, A., Tchuenche, J.M., 2011. Mathematical analysis of a cholera model with public health interventions. Biosystems 105, 190–200.

National Hurricane Center, 2020. Tropical cyclone climatology. URL: https://www.nhc.noaa.gov/climo/.

Neilan, R.L.M., Schaefer, E., Gaff, H., Fister, K.R., Lenhart, S., 2010. Modeling optimal intervention strategies for cholera. Bulletin of mathematical biology 72, 2004–2018.

OCHA, 2010a. Haiti cholera situation report 2. URL: https://reliefweb.int/sites/reliefweb.int/files/resources/FB97D45BDA59D54FC12577C6002FC3D0-Full_report.pdf.

OCHA, 2010b. Haiti cholera situation report 3. URL: https://reliefweb.int/sites/reliefweb.int/files/resources/4165CCA36E6031A9492577C7001F4A7A-Full_Report.pdf.

OCHA Haiti, 2019. Location of health facilities in haiti. URL: https://data.humdata.org/dataset/location-of-health-facilities-hospital-in-haiti.

Olson, D., Fesselet, J., Grouzard, V., 2018. Management of a cholera epidemic. Switzerland: MédecinsSansFrontières .

Olstad, M.W., Verås, H.G., 2019. Data-driven predictive modeling for production optimization .

Pan American Health Organization, 2010a. Paho position on cholera vaccination in haiti. URL: https://www.paho.org/hq/dmdocuments/2010/PAHO_position_cholera_vaccination.pdf.

Pan American Health Organization, 2010b. Paho responds to cholera outbreak in haiti. URL: https://reliefweb.int/report/haiti/paho-responds-cholera-outbreak-haiti.

Pan American Health Organization, 2014. Haiti to launch cholera vaccination with paho/who support. URL: https://www.paho.org/hq/index.php?option=com_content&view=article&id=9788:2014-haiti-to-launch-cholera-vaccination-with-pahowho-support&Itemid=135&lang=en.

Pan American Health Organization, 2017. Salud en las americas, haiti. URL: https://www.paho.org/salud-en-las-americas-2017/?p=4110.

Pan American Health Organization, 2020a. Haiti reaches one-year

free of cholera. URL: https://www.paho.org/en/news/23-1-2020-haiti-reaches-one-year-free-cholera.

Pan American Health Organization, 2020b. Paho eoc situation report - cholera outbreak in haiti. URL: https://reliefweb.int/updates?advanced-search=%28D6819_S1980%29.

Pinsky, M., Karlin, S., 2010. An introduction to stochastic modeling. 4 ed., Academic press.

Powell, W.B., 2007. Approximate Dynamic Programming: Solving the curses of dimensionality. volume 703. John Wiley & Sons.

Rachaniotis, N.P., Dasaklis, T.K., Pappis, C.P., 2012. A deterministic resource scheduling model in epidemic control: A case study. European Journal of Operational Research 216, 225–231.

Rahaman, M.M., Majid, M., Alam, A.J., Islam, M.R., 1976. Effects of doxycycline in actively purging cholera patients: a double-blind clinical trial. Antimicrobial agents and chemotherapy 10, 610–612.

Range, T.M., 2019. Lecture notes in sequential decision making.

Rawls, C.G., Turnquist, M.A., 2010. Pre-positioning of emergency supplies for disaster response. Transportation research part B: Methodological 44, 521–534.

Rawls, C.G., Turnquist, M.A., 2012. Pre-positioning and dynamic delivery planning for short-term response following a natural disaster. Socio-Economic Planning Sciences 46, 46–54.

Ren, Y., Ordóñez, F., Wu, S., 2013. Optimal resource allocation response to a smallpox outbreak. Computers & Industrial Engineering 66, 325–337.

Rottkemper, B., Fischer, K., Blecken, A., 2012. A transshipment model for distribution and inventory relocation under uncertainty in humanitarian operations. Socio-Economic Planning Sciences 46, 98–109.

Tanner, M.W., Ntaimo, L., 2010. Iis branch-and-cut for joint chance-constrained stochastic programs and application to optimal vaccine allocation. European Journal of Operational Research 207, 290–296.

Tanner, M.W., Sattenspiel, L., Ntaimo, L., 2008. Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. Mathematical biosciences 215, 144–151.

Tuite, A.R., Tien, J., Eisenberg, M., Earn, D.J., Ma, J., Fisman, D.N., 2011. Cholera

epidemic in haiti, 2010: using a transmission model to explain spatial spread of disease and identify optimal control interventions. Annals of internal medicine 154, 593–601.

United Nations, 2011. Report of the United Nations in Haiti 2010. URL: https://reliefweb.int/sites/reliefweb.int/files/resources/F9DE84C8F12B844B8525781B0053C3F6-Full_Report.pdf.

Wang, H., Wang, X., Zeng, A.Z., 2009. Optimal material distribution decisions based on epidemic diffusion rule and stochastic latent period for emergency rescue. International Journal of Mathematics in Operational Research 1, 76–96.

Wang, J., Modnak, C., 2011. Modeling cholera dynamics with controls. Canadian applied mathematics quarterly 19, 255–273.

World Health Organization, 2010. Cholera in haiti - update. URL: https://www.who.int/csr/don/2010_10_28/en/.

World Health Organization, 2013. Oral cholera vaccine stockpile for cholera emergency response. Geneva: International Coordinating Group on Vaccine Provision for Cholera .

World Health Organization, 2016. Infectious diseases. URL: https://www.who.int/topics/infectious_diseases/en/.

World Health Organization, 2017a. Malaria vaccines. URL: https://www.who.int/immunization/research/development/malaria/en/.

World Health Organization, 2017b. Weekly epidemiological record, 2017, vol. 92, 34. Weekly Epidemiological Record= Relevé épidémiologique hebdomadaire 92, 477–500.

World Health Organization, 2017c. Weekly epidemiological record, 2017, vol. 92, 35 [full issue]. Weekly Epidemiological Record= Relevé épidémiologique hebdomadaire 92, 521–536.

World Health Organization, 2018. Managing epidemics: key facts about major deadly diseases. World Health Organization.

World Health Organization, 2019a. Cholera. URL: https://www.who.int/news-room/fact-sheets/detail/cholera.

World Health Organization, 2019b. Poliomyelitis. URL: https://www.who.int/news-room/fact-sheets/detail/poliomyelitis.

World Health Organization, 2019c. Vaccines. URL: https://www.who.int/topics/vaccines/en/.

World Health Organization, 2019d. WHO commemorates the 40th anniversary of smallpox eradication. World Health Organization URL: https://www.who.int/news-room/detail/13-12-2019-who-commemorates-the-40th-anniversary-of-smallpox-erad

Yaesoubi, R., Cohen, T., 2011. Dynamic health policies for controlling the spread of emerging infections: influenza as an example. PloS one 6.

Yarmand, H., Ivy, J.S., Denton, B., Lloyd, A.L., 2014. Optimal two-phase vaccine allocation to geographically different regions under uncertainty. European Journal of Operational Research 233, 208–219.

Zaric, G.S., Brandeau, M.L., 2001. Resource allocation for epidemic control over short time horizons. Mathematical Biosciences 171, 33–58.

Zaric, G.S., Brandeau, M.L., 2002. Dynamic resource allocation for epidemic control in multiple populations. Mathematical Medicine and Biology 19, 235–255.

Appendix A

Implementation Structure

This appendix illustrates the file hierarchy for the implementation of the thesis' solution methods. The implementation is attached as a .zip-file, in addition to the written thesis. To run the code, see the README.md document.

```
tio4905
├── README.md
├── data
├── figures
└── models
    ├── cholera_model
    │   ├── Case.py
    │   └── Region.py
    ├── resource_allocation_model
    │   ├── MarkovDecisionProcess.py
    │   ├── State.py
    │   └── ValueApproximationFunction.py
    ├── trained_vfa_models
    ├── Instance.py
    └── main.py
```

Appendix B

Resource Allocation Model

Table B.1: Sets used in the resource allocation model.

Set	Definition
$\mathcal{T}$	Set of time periods
$\mathcal{I}$	Set of regions
$\mathcal{M}$	Set of intervention types
$\mathcal{N}$	Set of facility types

Table B.2: Indices used in the resource allocation model.

Index	Definition
t	Time period $t \in \mathcal{T}$
i	Region $i \in \mathcal{I}$
m	Intervention type $m \in \mathcal{M}$
n	Facility type $n \in \mathcal{N}$

Table B.3: Parameters used in the resource allocation model.

Parameter	Definition
R_{tm}	Number of resources of type m available at time t
D_{tim}	Demand for resource m in region i at time t
B_{nm}	Stock capacity of intervention type m for facility type n
U_m	Deployment and utilization capacity per available personnel per time period for intervention type m
Q_n	Minimum number of personnel required to establish one facility of type n
L_{in}	Number of available locations for facilities of type n in region i
P	Number of available medical personnel

Table B.4: Variables used in the resource allocation model.

Variable	Definition
x_{tim}	Number of resources of type m allocation to region i at time t
y_{tin}	Number of facilities of type n established in region i at time t
z_{ti}	Number of medical personnel allocated to region i at time t

Figure B.1: Resource allocation model solved each time period t.

$$\min_{x_t \in \mathcal{X}_t} \mathbb{E}\left\{ \sum_{t \in \mathcal{T}} C_t^E(S_t, x_t, W_{t+1}) \right\} \tag{B.1}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{I}} x_{tim} \leq R_{tm} \qquad \forall m \in \mathcal{M} \tag{B.2}$$

$$x_{tim} \leq D_{tim} \qquad \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{B.3}$$

$$x_{tim} \leq \sum_{n \in \mathcal{N}} B_{nm} y_{tin} \qquad \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{B.4}$$

$$\sum_{m \in \mathcal{M}} U_m x_{tim} \leq z_{ti} \qquad \forall i \in \mathcal{I} \tag{B.5}$$

$$\sum_{n \in \mathcal{N}} Q_n y_{tin} \leq z_{ti} \qquad \forall i \in \mathcal{I} \tag{B.6}$$

$$y_{tin} \leq L_{in} \qquad \forall i \in \mathcal{I}, \forall n \in \mathcal{N} \tag{B.7}$$

$$\sum_{i \in \mathcal{I}} z_{ti} \leq P \tag{B.8}$$

$$x_{tim} \in \mathbb{Z}^+ \qquad \forall i \in \mathcal{I}, \forall m \in \mathcal{M} \tag{B.9}$$

$$y_{tin} \in \mathbb{Z}^+ \qquad \forall i \in \mathcal{I}, \forall n \in \mathcal{N} \tag{B.10}$$

$$z_{ti} \in \mathbb{Z}^+ \qquad \forall i \in \mathcal{I} \tag{B.11}$$

Appendix C

Solution Method Procedures

This appendix contains pseudocode for the procedures not included in Chapter 7 for readability purposes. The pseudocode is not meant as an exact description of the coded algorithm, but rather describing the key aspects necessary to understand the solution methods.

Algorithm 7: TRAININGDECISION(S_t, ϵ)

Input : State S_t,
Random decision probability ϵ.
Output: Feasible decision x_t.

```
Draw random number, ξ1 ~ U(0,1);
if ξ1 < ε then
    Get random decision, x_t ← RANDOMDECISION(S_t);
else
    Draw random number, ξ2 ~ U(0,1);
    if ξ2 < ε then
        Get greedy decision, x_t ← GREEDYPOLICYDECISION(S_t);
    else
        Get ADP decision, x_t ← LOCALSEARCHPOLICYDECISION(S_t);
return x_t;
```

Algorithm 8: RANDOMDECISION(S_t)

Input : State S_t.
Output: Feasible decision x_t.

Initialize infeasible solution, $x_t \leftarrow (+\infty)_{(i,m)\in\mathcal{I}\times\mathcal{M}}$;
while *not* ISFEASIBLE(x_t) **do**
    Draw random, feasible resources to be allocated,
     $availableResources \sim$ RANDOMINTEGER$(0, S_t.resources)$;
    Shuffle iteration order of regions, $\mathcal{J} \leftarrow$ SHUFFLE$(\mathcal{I})$;
    for *intervention type* $m \in \mathcal{M}$ **do**
        for $i = 1, \ldots, |\mathcal{J}| - 1$ **do**
            Allocate random amount of available resources,
             $x_{tim} \sim$ RANDOMINTEGER$(0, availableResources[m])$;
            Update available resources,
             $availableResources[m] \leftarrow availableResources[m] - x_{tim}$;
        Allocate remaining resources to last region,
         $x_{t|\mathcal{J}|m} \leftarrow availableResources[m]$;
return Random decision, x_t;

Algorithm 9: REALLOCATION($R, \mathcal{I}, r, S, \alpha$)

Input : Initial resource pool R,
Regions $\mathcal{I}$,
Previous initial allocation r,
Previous states S,
Transfer step-size α.
Output: Initial resource allocation to all regions.

for *each region* $i \in \mathcal{I}$ **do**
    for *each intervention type* $m \in \mathcal{M}$ **do**
        Calculate marginal benefit, $e_{im} \leftarrow \frac{r_{im}}{S_{\mathcal{T}i.fatalities}}$
Get best resource transfer indices, $(i, j, m) \leftarrow \arg\max_{i,j,m}(e_{im} - e_{jm})$;
Transfer resources from region i, $r_{im} \leftarrow r_{im} - \alpha$;
Transfer resources to region j, $r_{jm} \leftarrow r_{jm} + \alpha$;
return New resource allocation, $r = (r_{im})_{(i,m)\in\mathcal{I}\times\mathcal{M}}$;

Appendix D

Cumulative Costs

The following figures show the cumulative fatalities plots for all scenarios included in the computational study. All plots include all four policies investigated.

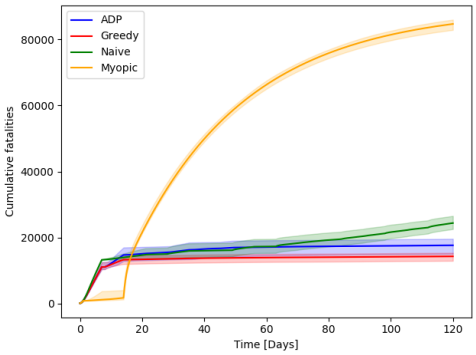

(a) Base case, calibrated epidemic model

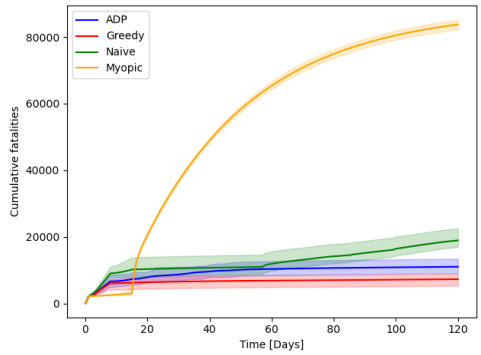

(b) Delayed, calibrated epidemic model

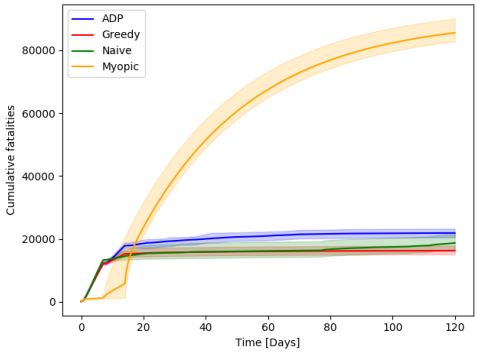

(c) Kit size 500, calibrated epidemic model

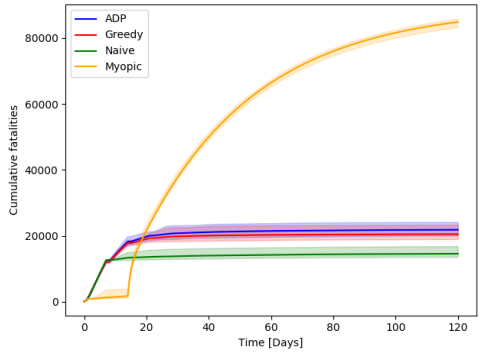

(d) Kit size 100, calibrated epidemic model

Figure D.1:

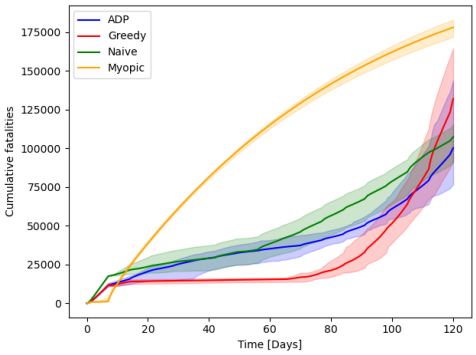

(a) Base case, alternative epidemic model

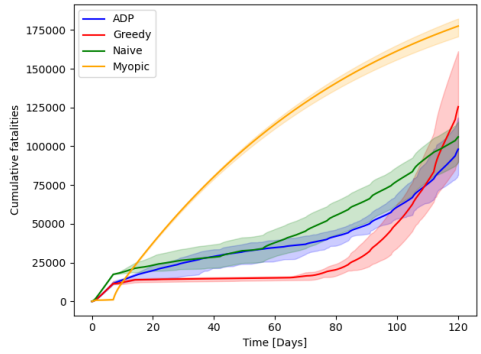

(b) One week vaccine arrival, alternative epidemic model

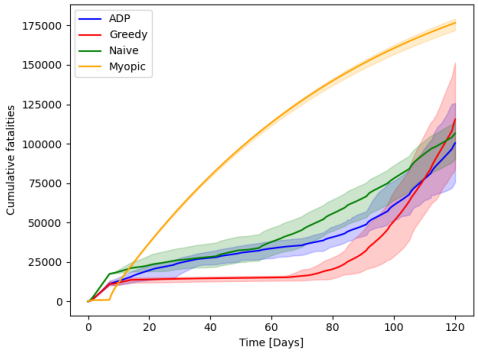

(c) Immediate vaccine arrival, alternative epidemic model

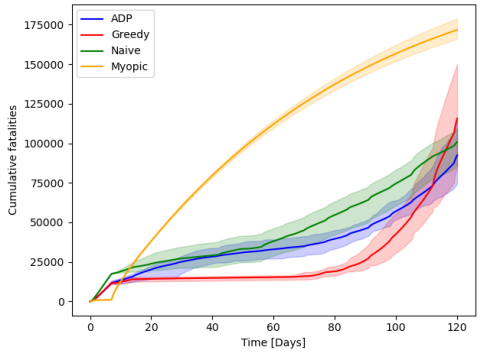

(d) 600 000 vaccines, alternative epidemic model

Figure D.2:

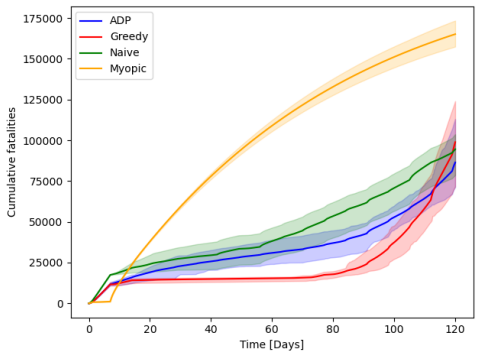

(a) 800 000 vaccines, alternative epidemic model

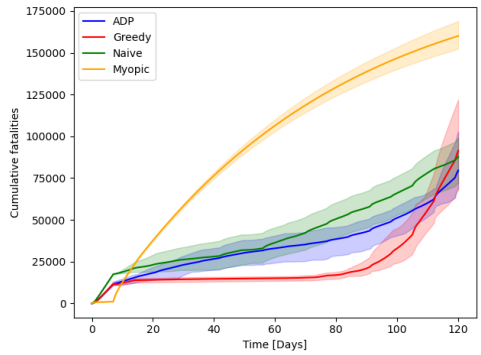

(b) 1 000 000 vaccines, alternative epidemic model

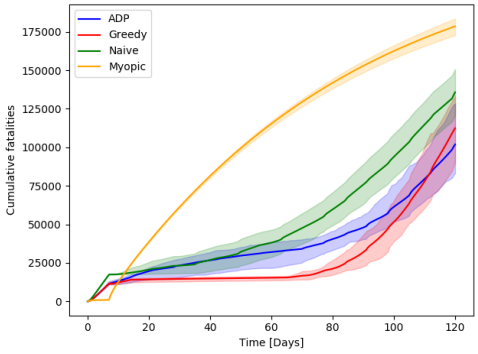

(c) 300 000 rehydration solutions, alternative epidemic model

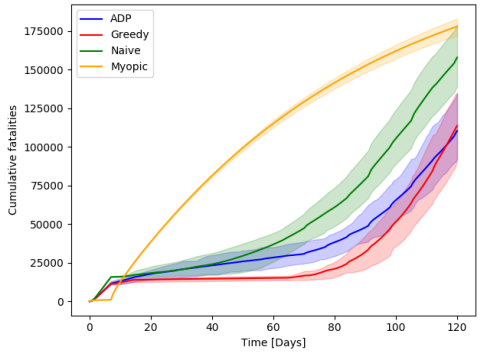

(d) 400 000 rehydration solutions, alternative epidemic model

Figure D.3:

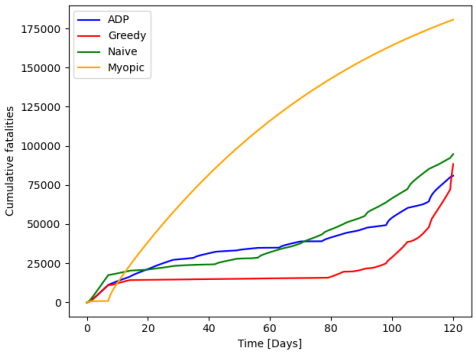

(a) Calibrated dispersal rate, alternative epidemic model

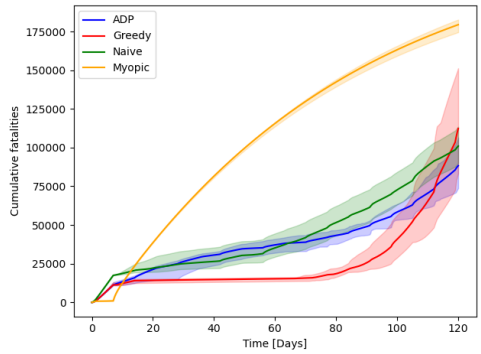

(b) Low variance dispersal distribution, alternative epidemic model

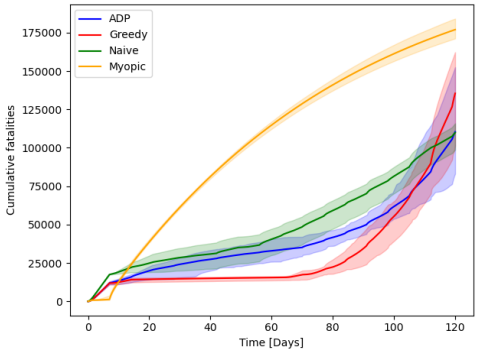

(c) High variance dispersal distribution, alternative epidemic model

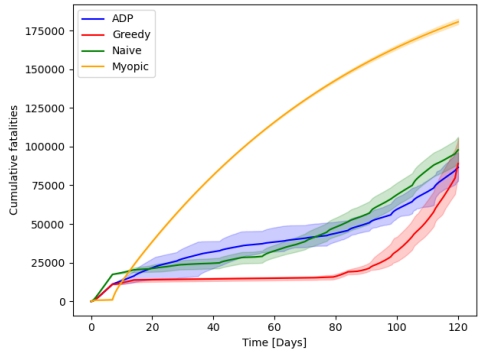

(d) Low maximum dispersal rate distribution, alternative epidemic model

Figure D.4:

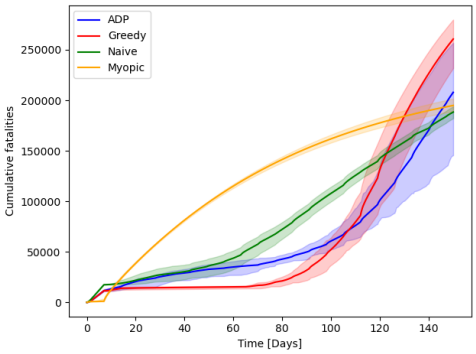

(a) 150 day horizon, alternative epidemic model