Uncertainty-Driven Policies for Resource Allocation in Epidemics Response

E. J. L. den Brok

Master Thesis for the MSc. Engineering and Policy Analysis at Delft University of Technology

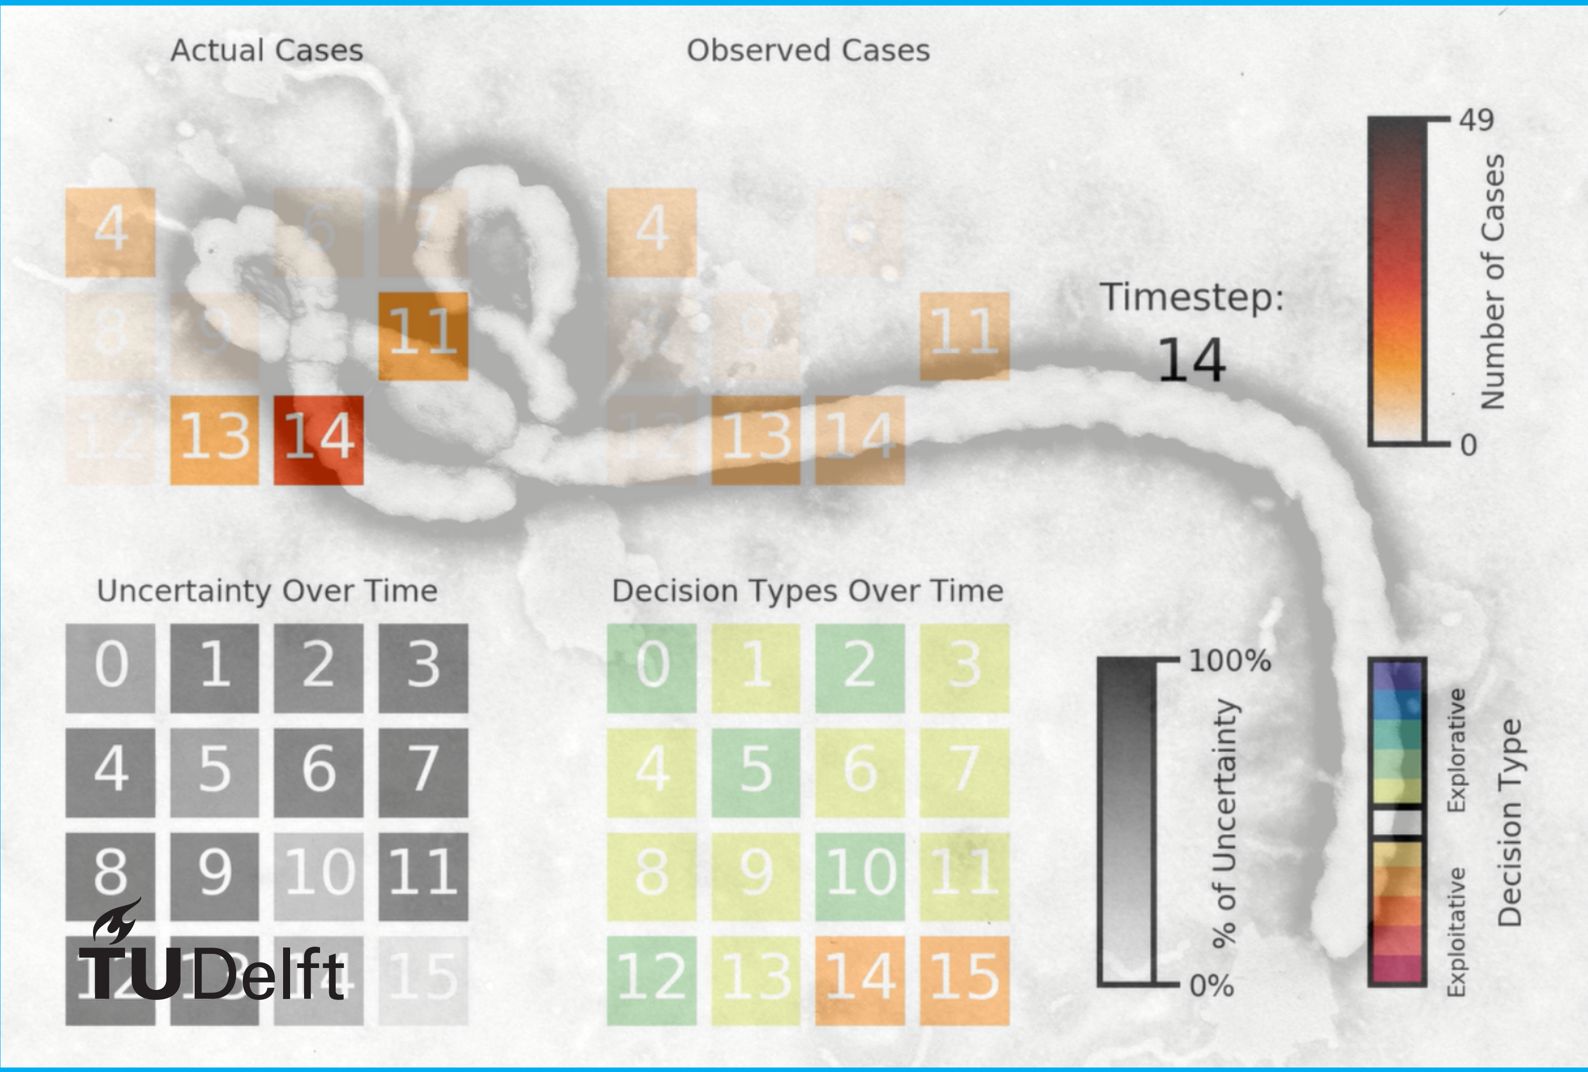

TUDelft

Uncertainty-Driven Policies for Resource Allocation in Epidemics Response

by

E. J. L. den Brok

to obtain the degree of Master of Science
at Delft University of Technology,
to be defended publicly on Monday August 19, 2019 at 3:00 PM.

Student number:	4738187	
Project duration:	Febuary 2019 – August 2019	
Thesis committee:	Dr. T. C. Comes	TU Delft, First supervisor
	Dr. ir. J. H. Kwakkel	TU Delft, Second supervisor
	Prof. dr. B. A. Van de Walle	TU Delft, Chair

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Preface

Dear reader,

In front of you lie six months of hard work, but they are the result of a journey I started more than a year ago, when I first joined the HumTechLab of TU Delft. After getting a taste of all the different ways in which the work done at TU Delft could contribute to the humanitarian world, I dove into the world of epidemics, humanitarian response and logistics, and simulation modelling, in the hope of making a valuable contribution to the invaluable work that humanitarian organisations carry out worldwide. The end result is found in this thesis.

Though writing a thesis sometimes feels like a lonely affair, I certainly could not have done it alone, and I have a number of people I want to thank here. First and foremost my graduation committee: Tina, thank you for all your encouragement, your ideas, and for always asking the right critical questions that helped me move forward and improve my work. Jan, thank you for your pointers in the deep uncertainty field, helping me with the implementation of my model, and of course for developing the EMA workbench! Bartel, thank you for keeping me - and everyone in the HumTechLab - motivated. Whether it was in sharing your experiences or your contacts, you ensured that I was always aware of the reality humanitarians operate in and how we could contribute to that.

To everyone else in the HumTechLab: I had a great time working alongside all of you, and being able to chat about what we were all doing over cookies and tea made it all the more fun. Kenny, thank you for giving us your coffee card - you are a true humanitarian.

To my friends and family: thank you for your support, whether it came in the form of a good conversation, a few drinks, a dinner cooked with love, cute pictures of animals, or simply by voicing your belief in me. It really helped, and I really appreciate it!

I would also like to express my gratitude to all the people working in the humanitarian sector who I spoke to for my thesis. I was blown away by how generous you all were with sharing your time, experiences, and ideas with me. I learnt a lot from talking to you, and it made me extremely motivated to do the best work I could. I hope that this thesis can make a small contribution to the amazing work you do all over the world.

Emma den Brok
Delft, August 2019

Executive Summary

Research Context

Humanitarians and global health actors come to the aid of many people every year, with the aim of preventing disease, increasing wellbeing, and providing (medical) aid to those suffering from disease. One of the contexts in which they operate is that of an epidemic. An epidemic is dynamic by nature and provides a complex and evolving environment in which medical aid needs to be provided. A key aspect in a response to an epidemic is logistics – specifically the allocation of resources such as personnel and medical supplies. These resources are often limited, calling for a targeted and strategic response.

There is a variety of studies tackling the problem of resource allocation in the context of an epidemic, which include sequential decisions as the epidemic evolves, as well as the choice between several locations to which resources can be sent. However, these studies often assume decision-makers have complete information on the situation at hand and can make “perfect” choices. In reality, due to the large number of actors involved in a response, poor (telecommunication) infrastructure, and the fact that an epidemic is a moving target due to its dynamic nature, decision-makers often have to deal with incomplete and uncertain information on the number of patients and the way the epidemic is evolving.

One of the goals of this thesis is to incorporate the uncertainty decision-makers face into the decision problem. A key assumption connected to this goal is that the uncertainty as experienced by the decision-maker can be reduced by sequential decisions. That is, as the response evolves and the decision-maker places resources, they will obtain more knowledge on those regions in which they have established a presence. Additionally, this thesis also considers multiple objectives on which the success of a response is evaluated, reflecting the different perspective of stakeholders involved in an epidemic. The main research question of this thesis is formulated as follows:

> What is the value of explicitly incorporating uncertainty and its reduction by sequential treatment centre placement decisions for epidemics response?

Approach

The research question is approached with Simulation Modelling and Decision-Making under Deep Uncertainty methods. In order to develop a realistic model, the 2014 Ebola epidemic in West Africa is used as a case study, from which epidemiological and demographic data is used. A simulation model is developed which represents the spread of Ebola over 16 regions over a period of six months, approximating the situation in Sierra Leone from June to December 2014. The decision-maker has no or uncertain knowledge on the state of the epidemic in each of the 16 regions. Every week, the decision-maker can choose to take either an explorative or an exploitative action. Explorative actions are aimed at reducing uncertainty by sending resources to a region of which the decision-maker has no or highly uncertain knowledge. Exploitative decisions are aimed at using resources to treat known cases.

In order to identify the policies that prescribe what type of action a decision-maker should take, Direct Policy Search (DPS) is used. DPS uses closed-loop policies: policies which base decisions on some observation of the current state in the system. In the context of this thesis, the policy is dependent on the level of uncertainty in the system as experienced by the decision-maker. The policy then dictates whether the decision-maker should take explorative or exploitative action.

In DPS, a Many-Objective Evolutionary Algorithm (MOEA) is used to identify optimal policies based on a reference scenario. The resulting policies are then tested on robustness by evaluating their performance on a large number of different scenarios, which are created by varying model parameters. The model parameters which are used to generate scenarios in this thesis are the transmission rate, the initial number of cases (infections) at the start of the simulation, and the travel rate. The policies are evaluated on five objectives: effectiveness, efficiency, equity in met demand, equity in arrival time, and time until the epidemic is contained.

Findings

Analysis on the performance of different policies shows that explorative policies (policies which incorporate a large amount of explorative action) avoid worst-case scenarios in which an epidemic would be discovered too late or not at all by purely exploitative policies. This effect is most visible in scenarios with epidemics that evolve slowly or if isolated regions are strongly affected. However, in scenarios where exploitative policies are successful, explorative policies fail to reach the same level of performance as they constantly trade-off meeting demand against spending resources on reducing uncertainty. This also leads to poorer performance of explorative policies in terms of efficiency and time until containment. Since explorative policies provide the decision-maker with a more complete situational awareness, they are also valuable when considering the equity objectives. Explorative policies perform better in terms of equity in arrival time because they spread the placement of resources over all regions. However, for equity in met demand the benefit of situational awareness is negated by the trade-off between explorative and exploitative actions, as resources are not solely focussed on meeting demand.

The main contributions of this thesis are two-fold: From the humanitarian perspective, the results provide important insight into the risks and benefits of current practice (which is assumed to be fully exploitative), and the value which other strategies for resource allocation during epidemics can have for the sector. The other contribution is in the field of deep uncertainty research. By incorporating uncertainty reduction through sequential decisions, and by linking the experienced level of uncertainty directly to the policy function, this thesis shows the value of considering uncertainty within the simulation model itself. This adds a new dimension to the method of Direct Policy Search.

Contents

1 Introduction **1**
1.1 Epidemics - Complex and Dynamic Disasters . . . 1
1.2 Problem Statement . . . 2
1.3 Literature Review . . . 2
1.3.1 Humanitarian Logistics . . . 2
1.3.2 Resource Allocation in Epidemics Response . . . 2
1.4 Research Gap and Research Question . . . 5

2 Research Approach **7**
2.1 Deep Uncertainty . . . 7
2.2 Sub-questions . . . 8
2.3 Methods . . . 8
2.3.1 Conceptualizing the System . . . 8
2.3.2 Approaching Uncertainty as an Exploration Vs. Exploitation Problem . . . 8
2.3.3 Policy Search . . . 9
2.3.4 Simulation Model & Case Study . . . 9
2.4 Structure of this Thesis . . . 10

3 Model Conceptualisation **11**
3.1 Delimitation of the Problem . . . 11
3.2 Problem Owner and Objectives . . . 11
3.2.1 Actors involved in the Ebola Response . . . 11
3.2.2 Problem Owner . . . 12
3.2.3 Objectives . . . 14
3.3 Basic Model Structure . . . 15
3.3.1 Compartmental Model . . . 15
3.3.2 Decision Levers . . . 15
3.4 Uncertain Factors . . . 17
3.5 Conclusion . . . 17

4 Uncertainty Reduction **19**
4.1 Conceptualisation of Uncertainty Reduction . . . 19
4.1.1 Number of Infected Individuals . . . 20
4.1.2 Transmission Rate . . . 21
4.1.3 Travelling Behaviour . . . 21
4.1.4 Surveillance Teams . . . 22
4.2 Information Delay . . . 22
4.3 Explorative vs. Exploitative Decisions . . . 23
4.3.1 Defining explorative and exploitative decision . . . 23
4.3.2 Choosing the type of decision . . . 24
4.4 Conclusion . . . 24

5 Model Design **25**
5.1 Model Requirements . . . 25
5.2 Assumptions . . . 25
5.3 General Structure . . . 28
5.3.1 Compartmental Model . . . 28
5.3.2 Implementation of Uncertainty . . . 30
5.3.3 Decision-making Module . . . 31
5.3.4 Objectives . . . 31
5.3.5 Implementation . . . 32

6 Model Validation **33**
6.1 Validity of the Simulation Model . . . 33
6.1.1 Epidemiological Model . . . 33
6.1.2 Model Outcomes . . . 34
6.2 Validity of the Simulation Model for the Research Question . . . 34
6.3 Validity of the Research approach . . . 35

7 Experiments **37**
7.1 EMA workbench . . . 37
7.1.1 Exploratory Analysis . . . 37
7.2 Base Scenario . . . 39
7.3 Direct Policy Search using MOEA . . . 39
7.4 Robustness Testing . . . 40

8 Results **43**
8.1 Model Behaviour . . . 43
8.1.1 Outcome Distributions . . . 43
8.1.2 Influence of Uncertain Factors . . . 45
8.1.3 Runtime Behaviour . . . 49
8.2 Policy Performance. . . . 56
8.2.1 Policy Behaviour per Classification . . . 56
8.2.2 Runtime Behaviour . . . 59
8.2.3 Robustness Scores. . . . 64

9 Discussion **67**
9.1 Uncertain Factors. . . . 67
9.2 Behaviour and Performance of Policies . . . 68
9.2.1 Exploitative Policies . . . 68
9.2.2 Explorative Policies. . . . 69
9.2.3 Mixed Policy . . . 69
9.2.4 Value of the Different Policy Types . . . 70
9.3 Robustness of the Policies . . . 70
9.4 Limitations . . . 71
9.5 Conclusion . . . 72

10 Conclusion **75**
10.1 Summary . . . 75
10.2 The value of explicitly incorporating uncertainty in resource allocation strategies . . . 76
10.3 Scientific Contribution . . . 76
10.4 Societal Relevance . . . 77
10.5 Future Research . . . 77

References **79**

A List of Abbreviations used in this Thesis **85**

B Uncertainty Reduction - Lookup Tables and Functions used in the Simulation Model **87**

C Simulation Model Objectives **89**

D Simulation Model Parametrization **91**

E Validation of the Ensemble Size **93**

F MOEA Outcomes **95**

G Scatterplots of All Policies **99**

H Runtime Behaviour **107**
H.1 Exploitation-390 Policy . . . 107
H.2 Exploration-185 Policy . . . 107

I Code - Overview **111**

J Interviews **113**

1
Introduction

1.1. Epidemics - Complex and Dynamic Disasters

Every year, millions of people face the reality of disasters, natural or man-made. In 2017 alone, over 11000 people worldwide lost their life or went missing, while over a million became homeless as a result of 183 natural disasters and 118 man-made disasters (Bevere, Schwartz, Sharan, & Zimmerli, 2018).

Epidemics are disasters that cause large amounts of suffering and potentially many deaths. Recent examples of outbreaks include the cholera outbreak in Haiti after the earthquake in 2010, which resulted in 470,000 known cases and over 6000 deaths (Centers for Disease Control and Prevention, 2011) and the 2014 Ebola epidemic in West Africa, which infected over 28,000 people and killed over 11,000 (Kaner & Schaack, 2016). These examples also illustrate that epidemics occur in many contexts: spontaneously, or in the aftermath of another disaster. Similarly, epidemics can evolve rapidly or take place over many years, as is the case for HIV.

Epidemics will remain a challenge in the 21st century. According to the UN (United Nations, 2017), the world population is expected to have grown to 11 billion by 2100 with the least developed countries accounting for a large part of this growth. Population growth, along with urbanisation and climate change increase the risk of infectious disease and epidemics (Alirol, Getaz, Stoll, Chappuis, & Loutan, 2011; Gholipour, 2013; World Health Organization, 2003).

Those who respond to (the threat of) an epidemic have to operate in complex environments with multiple actors, who can have conflicting objectives, and face uncertainties such as the size and location of demand for medical aid and supplies (Liberatore, Pizarro, de Blas, Ortuño, & Vitoriano, 2013; van der Laan, van Dalen, Rohrmoser, & Simpson, 2016). Problems that arise from response in under such conditions are exemplified by the 2014 Ebola epidemic. As no cure or vaccine was available at the time, an important part of the response was focussed on quarantining and providing basic care to those infected with Ebola Virus Disease (EVD) [1]. Yet coordination between governments of the three affected countries, the WHO, and NGOs was poor, and aid workers could not even get basic information such as the number of Ebola Treatment Centres (ETC) (Gettleman, 2014). Often, ETCs were over capacity as soon as they were opened due to the large influx of (suspected) patients (Payne, 2014). The struggle to keep up with the spread of the epidemic was voiced by an MSF coordinator who said *“Everything we do is too small and too late […] We’re always running after the epidemic.”* (Diallo & Dilorenzo, 2014). Later in the response, the WHO claimed that on a national level enough beds were available, but admitted that on a district level some ETCs were still over capacity, while others had spare beds (World Health Organization, 2014c). This shows that even once enough resources were available, their proper distribution in the face of an uncertain and dynamic environment remained a challenge.

[1] A full list of all abbreviations used in this thesis can be found in Appendix A.

1.2. Problem Statement

In an epidemics context, decision-makers face the difficulty of choosing where to send resources in a constantly evolving environment. Additionally, due to a lack of coordination and communication, or simply due to exogenous factors (such as the availability of telecommunication technologies) information necessary to make these decisions can be missing or uncertain. Moreover, multiple different actors (i.e. governments, NGOs, international community actors) can be involved in a response, each with a different perspective and with different goals.

There is a need for decision-support in such situations, which takes into account the (un)availability of information, conflicting objectives, and the fact that as an epidemic spreads, it essentially becomes a “moving target.” This thesis aims to provide this decision-support by investigating resource allocation during epidemics in the form of a facility placement problem. It will consider how information uncertainty influences the performance of resource allocation policies, how decisions can be used strategically to reduce uncertainty, and evaluate if such policies lead to better allocation of resources. To reflect the multi-actor context in which an epidemic takes place, multiple objectives will be considered.

In order to gain a better understanding of the subject and to become familiar with the work already done in this area, the next section reviews academic literature on humanitarian logistics, and resource allocation specifically in the context of an epidemic.

1.3. Literature Review

1.3.1. Humanitarian Logistics

Two types of humanitarian disasters can be distinguished: slow onset disasters (such as an epidemic or famine) and sudden-onset disasters (i.e. a flood or earthquacke) (Charles & Lauras, 2011). Additionally, humanitarian operations are categorised by four phases: mitigation, preparedness, response, and recovery (Çelik et al., 2012). Within humanitarian response, humanitarian logistics is an area of research and a branch of operations that deals with the supply, transport and allocation of goods, people and information in disaster response. Though humanitarian logistics play a vital role in a response, it only started to develop as an area of research in the early 2000s (Balcik & Beamon, 2008). Humanitarian logistics differs from its commercial counterpart as its focus is on meeting the needs of the end-receiver, rather than focussing on meeting the objectives of the supplier (Çelik et al., 2012). When comparing humanitarian facility location problems with comparable commercial issues, specific differences are that in the humanitarian context demand can spike suddenly and is highly unpredictable in terms of time and location. Additionally, the stakes associated with decision-making are high and resources are limited (Balcik & Beamon, 2008) (Holguín-Veras, Jaller, Van Wassenhove, Pérez, & Wachtendorf, 2012).

Dasaklis, Pappis, and Rachaniotis (2012) review humanitarian logistics for epidemics control. They describe the role of logistics in epidemics response as dealing with resource allocation, cold supply-chain setup (for vaccines and blood samples) and managing the availability of staff en resources. The problem studied in this thesis is a resource allocation problem during the response phase of an epidemic. A literature review on studies considering the same problem is carried out in the next section.

1.3.2. Resource Allocation in Epidemics Response

Ahmadi-Javid, Seyedi, and Syam (2017) review the work done on healthcare facility location problems, which also includes non-emergency contexts. However, they identify no research on the emergency placement of medical facilities in an epidemics context; as all four studies they place in this category are in the context of a sudden-onset natural disaster. They do identify a study by Murali, Ordóñez, and Dessouky (2012) which considers the placement of points-of-dispersion for vaccines after a bio-terrorist attacks. This study will be discussed in more detail below.

Since there are few studies specifically considering medical facility placement in an epidemics response, the literature review will widen its scope to include studies which model the allocation of medical resources during epidemics. These studies will be evaluated on their incorporation of four characteristics, which are partly drawn from the classification suggestions by Ahmadi-Javid et al. (2017):

- Dynamic Epidemiological Model - Does the research incorporate a (compartmental) disease model that evolves over time and therefore requires sequential decision-making?
- Spatial Resource Allocation - Does the research consider a spatial allocation problem, i.e. with multiple regions or demand points (in contrast with allocation aimed at specific demographics).
- Evaluation Criteria - On how many, and which criteria the performance of the model is optimized.
- Uncertainty - Whether uncertainties in parameters or mechanisms (i.e. model structure) are considered systematically.

In order to create a state-of-the-art overview and show the most recent work, only studies published after 2010 were considered. Twelve relevant studies were identified. The evalutation of all these studies on their inclusion of the defined characteristics is shown in Table 1.1. Most studies frame the resource allocation problem as an optimization problem which can be solved mathematically. Mixed-integer and linear programming approaches are often used as a method for this.

Mbah and Gilligan (2011) study how limited resources can be assigned to several regions, each of which is represented by a SIRS-model (Susceptible-Infected-Recovered-Susceptible, explained in more detail in Chapter 3.) which evolves over time, meaning their problem is dynamic. They use mathematical analysis to identify the conditions for optimal solutions and then use numerical simulation to test strategies with the goal to minimize the (discounted) burden of disease. They also discuss the effect of the strategies on the trade-off between efficiency versus effectiveness. A static allocation problem is considered by Murali et al. (2012), whose study was also identified in the literature review by Ahmadi-Javid et al. (2017). They use an integer programming model to solve a resource allocation problem framed as a coverage problem, in which geographically spread demand points need to be covered by supply points. They aim to maximize demand covered, and they incorporate uncertainty by drawing the size of demand at a demand point from a known probability distribution.

Anparasan and Lejeune (2017) also consider a static allocation problem, in the context of the cholera outbreak that occurred after the 2011 earthquake in Haiti. They use integer linear programming to optimize the allocation of treatment facilities as well as ambulances over several regions with the goal to maximize the number of patients receiving treatment.

All other studies consider dynamic allocation problems with compartmental disease models as in Mbah and Gilligan (2011). Rachaniotis, Dasaklis, and Pappis (2012) use compartmental models to represent sub-populations to which limited medical resources must be scheduled. They aim to minimize the total number of infections and compare the performance of their optimized strategy to that of a random one, and incorporate uncertainty by evaluating both strategies over 1000 scenarios with randomly sampled parameters. The same authors also study the distribution of vaccines over sub-populations which are represented in a network structure in (Dasaklis, Rachaniotis, & Pappis, 2017). Again they incorporate uncertainty in the form of different scenarios by varying model parameters.

Other studies that incorporate scenarios are those by He and Liu (2015), Wanying, Alain, and Angel (2016) and Ren, Ordóñez, and Wu (2013). He and Liu (2015) use a compartmental model to forcast demand at resource centres and aim to minimize suffering by assigning resources to centres according to the forcast demand. Wanying et al. (2016) use linear programming to optimize the distribution of antibiotics in the case on an anthrax attack. Ren et al. (2013) uses mixed integer programming optmization to assign limited resources to several regions in the event of a bio-terrorist smallpox attack. In all three studies, a handful of different pre-determined scenarios is considered. This is not counted as incorporation of uncertainty, as the scenario's are presented as being "known" beforehand, and their small number means performance of strategies is not evaluated structurally.

The remaining studies do not incorporate uncertainty at all. Rachaniotis, Dasaklis, and Pappis (2017) revisit the resource scheduling problem first discussed in (Rachaniotis et al., 2012) and use a different solution strategy, but do not incorporate varying parameters. Liu and Zhang (2016) optimize the allocation of medicine to various regions over multiple timesteps with mixed integer programming, and use a compartmental model to forecast future demand each timestep. They aim to minimize cost

Article	Dynamic	Spatial	Criteria	Uncertainty
Mbah and Gilligan (2011)	x	x	3 (minimize burden of infection, efficiency, effectiveness)	
Murali et al. (2012)		x	1 (maximize demand covered)	
Rachaniotis et al. (2012)	x	x	1 (minimize total infections)	x
Ren et al. (2013)	x	x	1 (minimize deaths)	
He and Liu (2015)	x	x	1 (minimize suffering)	
Liu and Zhang (2016)	x	x	1 (minimize cost)	
Wanying et al. (2016)	x		1 (minimize deaths)	
Anparasan and Lejeune (2017)		x	1 (maximize demand covered)	
Auping et al. (2017)	x		n/a (no optimization)	x
Dasaklis et al. (2017)	x	x	1 (maximize demand covered)	
Rachaniotis et al. (2017)	x	x	1 (minimize new infections)	x
Büyüktahtakın et al. (2018)	x	x	1 (minimize deaths)	

Table 1.1: showing the categorization of the discussed literature according to the four formulated characteristics.

given a constraint on the minimum amount of demand that needs to be covered. Büyüktahtakın, des Bordes, and Kıbış (2018) also incorporate a multi-period planning horizon and aim to minimize the number of deaths by allocating resources in the context of the 2014 Ebola epidemic.

Though it does not involve a spatial component, and is based on different methods than the previously discussed research, a study by Auping, Pruyt, and Kwakkel (2017) is included here because of its explicit inclusion of uncertainty. Using system dynamics modelling to study the effect of different intervention methods and strategies on a population, they analyse the performance of different responses under many different conditions using scenario discovery.

Based on the literature review, the following observations can be made: Several studies exist that consider a spatial resource allocation problem with a dynamic disease model. Additionally, several studies consider a handful of scenarios and/or model structures to study the effect they have on the resource allocation optimization. Structural consideration of uncertainty, however, is not found in these studies, with the exception of Rachaniotis et al. (2012) and Rachaniotis et al. (2017) who evaluate policy performance by varying model parameters. Some studies consider uncertainty explicitly (i.e. Murali et al. (2012) who considers uncertainty in demand, and Auping et al. (2017) who perform scenario discovery) but none of these include both a dynamic disease model and a spatial component for the allocation problem. The number of criteria on which the resource allocation problem is optimized is almost always one: either to minimize cost (given a constraint on met demand) or to minimize deaths or suffering given a certain amount of resources. The only exception is Mbah and Gilligan (2011) who considers optimal strategies with regard to efficiency and effectiveness and the trade-off between the two.

1.4. Research Gap and Research Question

In the problem statement, it has been established that uncertainty on the spread and development of a disease poses a problem in epidemics response. However, from the literature review it is concluded that in resource allocation research for epidemics, uncertainty is not incorporated structurally. This may mean that the policies identified in these studies perform sub-optimally in reality or require conditions on information availability which cannot be met in a real-life response. It is therefore relevant to study the effect of uncertainty on allocation strategies.

Additionally, given the long time-frame over which epidemics take place, the decisions made in the response will be sequential. That is, not all resources will be placed at once. Instead, a decision made earlier in the response affects the state of the epidemic and therefore all decisions made after it. It is therefore also relevant to investigate how uncertainty evolves over time as the result of a placement decision. In the context of treatment placement decisions, it is hypothesized that placing resources such as a treatment centre reduces uncertainty in the spatial area surrounding it. It is important to take into account this interaction between action and uncertainty. It was studied in a master's thesis by Romijn (2018) in the context of a humanitarian response after a natural disaster. This thesis proposes to expand on this by including the level of uncertainty in the system as a variable used in decision-making.

Finally, current research does not take into account multiple objectives when evaluating proposed solutions. When considering resource allocation over several sub-populations (as is done in spatial resource allocation problems) in a humanitarian context, it is insufficient to only consider efficiency or cost. Concepts such as equity and impartiality should be incorporated in order to come to policies that are acceptable to humanitarian decision makers. Another argument for multiple objectives is the fact that a response is often carried out by multiple stakeholders, who might prioritize different objectives.

On the basis of the literature review and the identified research gap, the following research question is formulated:

> What is the value of explicitly incorporating uncertainty and its reduction by sequential treatment centre placement decisions for epidemics response?

The research approach that will be taken to answer this question, and the relevant sub-questions that need to be answered in order to formulate an answer to the main research question are the subject of the next chapter.

2

Research Approach

In the previous chapter, the problem of resource allocation during epidemics was introduced, and on the basis of a literature review the knowledge gap was identified. The knowledge gap showed that current research does not consider uncertainty systematically even though this is a major hurdle in actual responses. Additionally, it was shown that most relevant research only considers a limited number of objectives, which may not be reflective of the multi-actor effort that takes place during an epidemic. This knowledge gap lead to the formulation of a research question that aims to study the role of uncertainty in allocation problems and its interaction with allocation decisions:

> What is the value of explicitly incorporating uncertainty and its reduction by sequential treatment centre placement decisions for epidemics response?

The aim of this chapter is to define a a point of view and research approach with which this question will be approached, to define the sub-questions necessary to conduct the approach, and to provide the methods with which the sub-questions will be answered.

2.1. Deep Uncertainty

In its essence, the problem of allocation of resources during an epidemic is a problem in which a series of sequential decisions need to be made on where (limited) resources will be sent. The position taken in this thesis is that these sequential decisions are made under a certain level of uncertainty, which is in turn influenced by past decisions, and the question is whether incorporating this uncertainty explicitly in the resource allocation policy will lead to better strategies.

Decision Support under Deep Uncertainty is an area of research which aims to develop policies or strategies while acknowledging the uncertainty present in complex systems, which fits the problem statement of this thesis. The term “deep uncertainty” was first defined by Lempert (2003) to refer to situations in which decision-makers cannot agree on the structure of a conceptual model, on the probability distribution of the uncertain parameters and values, and on how to evaluate various possible outcomes. The three key approaches to deal with these uncertainties are exploratory modeling, adaptive planning, and decision-support (Haasnoot, Warren, & Kwakkel, 2019). Exploratory modelling uses simulation models to generate plausible future states of the system by varying uncertain parameters and/or model structures. By generating a large number (thousands) of scenarios, patterns in system behaviour and its relation to input values can be used to reason about the effect of uncertainties on the system. Adaptive planning means to design policies that can be changed over time, depending on how the system evolves. This means that how the policy changes is dependent on observed developments, rather than that it is planned out ahead of time. Decision support under uncertainty takes into account that in complex systems often multiple actors need to come to a decision. Therefore, decision support is aimed at showing the trade-offs between multiple objectives and presenting decision-makers with alternatives, rather than suggesting one predetermined plan of action. All together, these methods aim to find strategies or policies that are robust, meaning their performance is only impacted marginally by the uncertain future aspects.

It is evident that the deep uncertainty paradigm suits the problem of studying uncertainty in resource allocation during epidemics well. With the approaches described above in mind, the sub-questions are formulated in the next section.

2.2. Sub-questions

A few key components of the deep uncertainty approach are now clear: exploratory modelling and adaptive planning require a simulation model in order to gain insight into the future states of the model. In order to provide decision-support for real-life situations, the objectives used in this context also need to be identified. Additionally, since it is assumed uncertainty can be reduced as the result of a placement decision, the manner in which this happens needs to be developed. This results in the following sub-questions:

1. *What are the objectives on which decisions for treatment centre placement are optimized?*
2. *What are the main sources of uncertainty relevant for a treatment centre placement decision?*
3. *How can the reduction in uncertainty resulting from a treatment centre placement decision be conceptualized?*
4. *How can the sequential treatment centre placement decision problem and the decision-uncertainty reduction be represented in a simulation model?*
5. *Given the simulation model, what is the influence of system uncertainties on the performance of resource allocation policies?*
6. *Given the simulation model, which strategies for resource allocation decisions show robust performance?*

2.3. Methods

2.3.1. Conceptualizing the System

In order to answer any of the sub-questions, one first needs to have an understanding of the context in which an epidemics takes place, who is involved in the response, and what factors and decisions determine the evolution of an epidemic. Additionally, the system in which responders operate, and what information they have available to them (or becomes available to them over time) should be known. This is done by reviewing academic literature on the subject and by studying reports from humanitarians and the international community. In order to also get a first-person account of an epidemics response, and to validate assumptions that unavoidably need to made, this desk research will be supplemented by interviews with experts from the field (see Appendix J). With this knowledge and understanding, sub-questions 1 to 3 can be answered, providing a basis for the remainder of the research that is grounded in reality.

2.3.2. Approaching Uncertainty as an Exploration Vs. Exploitation Problem

Since sequential decisions need to be made in an uncertain environment, and the placement of a treatment centre can reduce uncertainty in parts (districts) of the environment, the problem can be seen as an exploitation vs. exploration problem (Memarzadeh & Pozzi, 2016), a conceptualisation that is also widely used for online decision problems in reinforcement learning.

An explorative decision would be to place a resource such as a treatment centre in a region with high uncertainty about the severity of the epidemic. There could be no patients at all in this region, or many. An exploitative decision would be to place a treatment centre in a region where there is certainty on the (minimum) number of patients. This does not mean uncertainty has to be low in these regions, but the lower bound of the current known range gives a guarantee on the number of patients.

The main research question considers the value of incorporating uncertainty. In order to understand what this value is, there needs to be a comparison between a policy that incorporates uncertainty and a policy that does not. A policy that does not incorporate uncertainty would be one that represents current practices, i.e. that is fully exploitative. A policy that also takes explorative actions can then be seen as one that incorporates uncertainty by aiming to reduce it.

2.3.3. Policy Search

In essence, a policy in a resource allocation problem would dictate for each point in the sequential decision process, whether an explorative or an exploitative action would be taken. Sub-question 6 then seeks to find such a policy that performs optimally (i.e. is robust) over a large scenario space.

It is proposed to use Direct Policy Search (DPS) to find such a policy. With direct policy search, a closed loop policy is optimized for multiple objectives on a simulation model using a multi-objective evolutionary algorithm (MOEA) (Giuliani, Castelletti, Pianosi, Mason, & Reed, 2015). A closed loop policy is dependent on one or more observable system variables. For the current problem, the system variable would be the overall level of uncertainty in the system as experienced by the decision maker. The benefit of this is that there is no fixed, optimal sequence of actions (or types of actions) that need to be taken based on the simulation model. Instead, the policy provides the optimal ratio of explorative versus exploitative action based on the current level of uncertainty in the system. This makes it easier to translate the results to a real-life decision-making context. Additionally, is not very computationally heavy, which allows for larger and more complex models (i.e. a high number of regions, to represent a large geographic area realistically). The DPS will yield a set of Pareto-optimal policies for the objectives identified in sub-question 1. These policies are optimized over a reference scenario.

Following Quinn, Reed, and Keller (2017) these policies can then be tested for robustness using exploratory modelling (i.e. by testing their performance over a large set of scenarios), thereby answering sub-question 6.

2.3.4. Simulation Model & Case Study

The DPS approach requires a simulation model, both for the policy search by the MOEA and for robustness testing. The simulation model will be developed as a system dynamics model. This allows for the spatial and temporal dynamics of the epidemic to be represented in a number of compartmental models, where each compartmental model corresponds to a geographic region. Given the availability of data, such a region will correspond to a district. The placement of resources in a district can then be represented by changing the parameters of the compartmental model representative of that district. Sub-question 4 therefore represents the implementation of the simulation model.

To create a simulation model, a case study is needed. The 2014 Ebola epidemic is chosen for several reasons. Firstly, given that main question is formulated as a facility location problem, it means that the placement of separate medical facilities is a necessary part of the response. This is true if a disease requires specific medical treatment (i.e. isolation) or if existing healthcare infrastructure is insufficient. Both were the case for Ebola. Secondly, for a system dynamics approach, delays between stocks (compartments) are necessary to represent flows. This means a disease needs to have an observable incubation period and period between symptom onset and death or recovery, which is the case for Ebola (WHO Ebola Response Team, 2014). Finally, there needs to be data available on various disease parameters as well as the effect of response measures in order to build a model. The Ebola epidemic has been well-documented. In fact, several compartmental models as well as models studying the effect of response measures have already been built (i.e. (Atkins et al., 2016; Büyüktahtakın et al., 2018; Kucharski et al., 2015)). This means the focus of the research can be on the decision-uncertainty interaction as an epidemiological model does not need to be constructed from scratch.

Given the simulation model, the effect of the uncertainties as identified in sub-question 2 on the policies can be studied using exploratory modelling to answer sub-question 5.

2.4. Structure of this Thesis

This thesis is structured as follows: Chapter 3 provides the model conceptualization by delimiting the problem at hand, identifying the problem owner and their objectives, and by setting up the basic model structure and identifying the uncertain factors in this structure. Therefore, sub-questions 1 and 2 are answered in Chapter 3. Sub-question 3 is the subject of Chapter 4, which conceptualises uncertainty reduction and defines explorative and exploitative decisions. The design of the actual simulation model is described in Chapter 5, thereby answering sub-question 4. The validity of the resulting model is discussed in Chapter 6. Chapter 7 provides the exact parametrization and experiments which are conducted with the simulation model. Chapter 8 and 9 then present the results of these experiments and their analysis and discussion, which allows for answers to be formulated to sub-questions 5 and 6. The final chapter, Chapter 10, provides a summary of the results, answers the main research question, considers the contributions of this thesis and possible directions for future research.

3
Model Conceptualisation

This chapter serves to identify and define the key elements of a model that represents both an epidemic and a facility placement problem under uncertainty. These key elements are: the delimitation of the problem, the problem owner and their objectives, the basic model structure, and the uncertain factors within the model.

Therefore, this chapter also serves to answer the first two research questions:

1. What are the objectives on which decisions for treatment centre placement are optimized?
2. What are the main sources of uncertainty relevant fo a treatment centre placement decision?

The structure of this chapter is as follows: First, the problem considered by the model is delimited. Secondly, in order to determine the problem owner, the actors involved in the Ebola response are discussed and one is chosen to be the problem owner. On this basis the objectives for the model are defined. Next, the basic model structure is outlined. Given the model structure, the uncertain factors within the model are determined.

3.1. Delimitation of the Problem

When responding to an Ebola outbreak, a broad scale of complementary measures are taken (i.e. case isolation, contact tracing, social outreach, and vaccination[1]). The proposed model will only consider decisions on the placement of ETCs (in terms of location and capacity) and surveillance teams, and therefore only considers the case isolation aspect of the response. The effect of other measures will be taken into account where relevant (i.e. contact tracing is an important factor in the reduction of uncertainty, as will be discussed in Chapter 4) but will not be considered as separate decision variables.

Additionally, an epidemics response is carried out by a multitude of actors (which will be discussed in detail in the section below). To model the interaction and cooperation between these actors is out of scope for this research (this is discussed in more detail in Chapter 6). Therefore, for the purpose of the model this process is not included and it is assumed there is one actor who makes the decisions on ETC or surveillance team placement. This actor is considered the problem owner and is the actor who requires the decision-support provided by the model. The problem owner is determined in the next section.

3.2. Problem Owner and Objectives

3.2.1. Actors involved in the Ebola Response

At the height of Ebola crisis, a large number of actors were involved - crisis response maps show that in some regions, more than 11 actors were active at one time (UNOCHA, 2014). These include govern-

[1]Vaccination trials were used in the latter stages of the West African Epidemic (Centers for Disease Control and Prevention, 2019) and are part of recent responses to Ebola in eastern DR Congo (World Health Organization, 2018).

mental actors (i.e. ministries of health), organisations concerned with public health (CDC, WHO), UN organizations (UNICEF, UNOCHA), and humanitarian organizations such as Save the Children, the International Federation of the Red Cross and Crescent (IFRC) and Medicines sans Frontières (MSF).

On September 19 2014, the UN established a special mission, UNMEER (UN Mission for Ebola Emergency Response) in order to scale up and coordinate the international response (United Nations, 2019). Representatives of affected countries, UN member states, international organisations such as the African Union and the African Development Bank Group, and humanitarian organisations involved in the response provided input or resources for the response. Within UNMEER, the WHO lead the development of a planned response (Secretariat of the World Health Organization, 2015). As part of the response from the international community, the African Union, the US, the UK, and France, amongst others, also deployed military resources to aid the response (Harman & Wenham, 2018).

A stylised representation of the actors involved in the response and their relationships is shown in Figure 3.1.

The actors involved in the response can be classified as global health actors and medical humanitarians as described by (Harman & Wenham, 2018). Global health actors are the WHO, ministries of health, and governmental bodies such as the CDC. Medical humanitarians are actors such as IFRC and MSF. This distinction is important, because global health actors operate differently from medical humanitarians and this impacts the objectives they have during an epidemics response. This will be elaborated on in Section 3.2.

To conclude, an epidemics response as seen for the Ebola epidemic involved a huge number of actors. In addition, the roles of and hierarchies between actors were often unclear (Harman & Wenham, 2018). Therefore, choosing one actor to be the problem owner and model-user is a simplification of reality. However, since the focus of this research is on modelling uncertainty and its interaction with decisions, such a simplification is necessary.

3.2.2. Problem Owner

Two actors appear as candidate problem owners: the WHO and UNMEER. Both actors had the mandate to coordinate the response to Ebola on an international level.

During the epidemic, the national ministries of health were in charge of the response at a national level, and it is important to recognize that in an international response, national governments still retain sovereignty over their country. Since the research considers an epidemic that crossed borders, and incorporating coordination between actors is out of scope for this research, the WHO and UNMEER are considered here instead of national governments.

The WHO is the *“directing and co-ordinating authority in international health work”*, and within its mandate it can support governments in times of emergency (World Health Organization, 2014a). The WHO monitored the response on 6 levels, including on the presence of and demand for Ebola treatment centres and referral centres (World Health Organization, 2014). In a strategy document published in August 2014, the WHO states it should *“coordinate international teams [...] and serve as a focal point for national and international teams”* (World Health Organization, 2014).

The other possible problem owner is UNMEER, which was created with the goal to improve the coordination of the response. However, UNMEER was a collection of organisations headed by the UN, and consists of many different actors with different resources and interests. Within UNMEER, WHO was the lead on case management and procurement of Ebola Treatment Centres (Secretariat of the World Health Organization, 2015). This means that within UNMEER, the WHO was the actor who would benefit from decision-support of the proposed model. Additionally, after UNMEER was disbanded on July 31 2015 the oversight of the UN response was transferred to the WHO.

Therefore, the WHO is assumed to be the problem owner. Choosing the WHO as the problem owner is not ideal, as the organization is still largely focussed on prevention and surveillance instead of response (Harman & Wenham, 2018). Additionally, the WHO’s reaction to the Ebola crisis has been widely critizised (see for example (Karlsen & Kruke, 2018; Médecins Sans Frontières, 2015; Moon et al.,

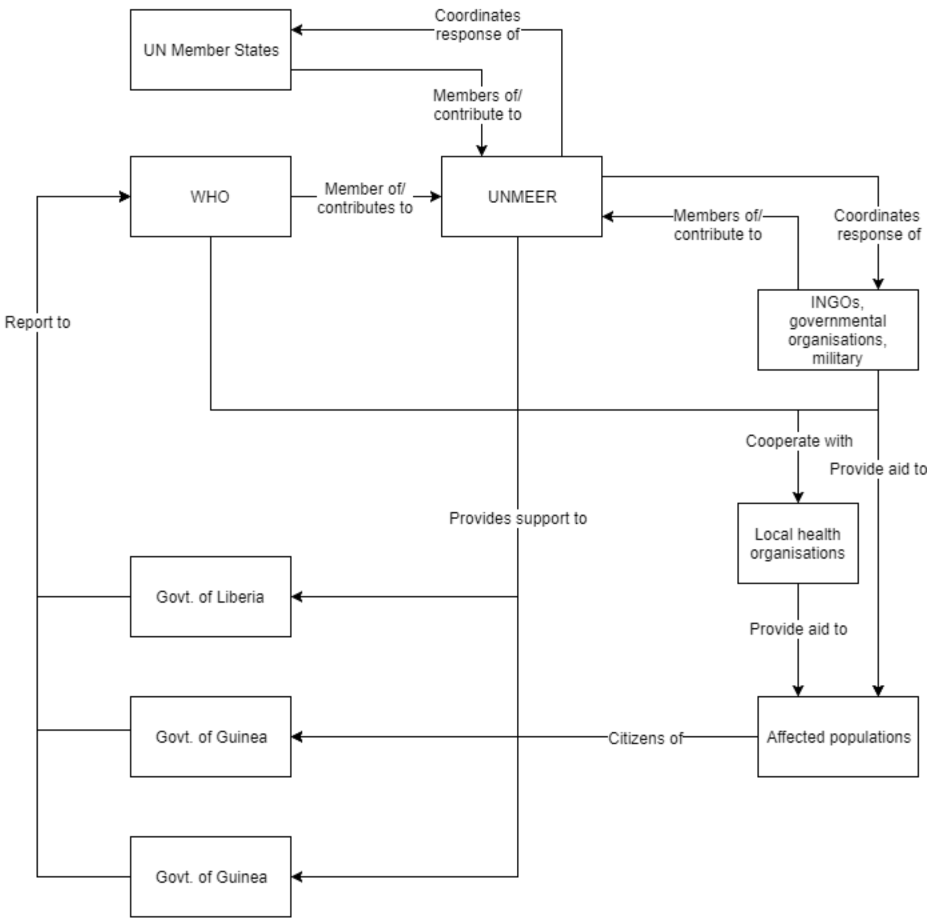

Figure 3.1: Formal Chart showing an stylised representation of the relationship between actors involved in the Ebola repsonse

2015)). Despite this, choosing the WHO as the problem owner is seen as a necessary and acceptable simplification in the context of this thesis.

3.2.3. Objectives

The WHO's mission statment is *"to promote health, keep the world safe and serve the vulnerable"* (World Health Organization, 2019). For the Ebola response, the WHO formulated its objective for countries dealing with high levels of infection as having full geographic coverage of Ebola response measures (World Health Organization, 2014b). Moreover, within UNMEER, a target was set by the WHO to have 70% of cases isolated by 1 December 2014, and 100% by 1 January 2015 (Secretariat of the World Health Organization, 2015). Translating this to model objectives, effectiveness is identified as the most important objective of the response. Effectiveness is defined as the extent to which response measures succeed in reducing the number of Ebola patients and the number of patients who die. Moreover, the WHO is dependent on its member states for funds and resources are not unlimited. Therefore, efficiency is also established as an objective. This is defined as minimizing the cost in dollars per life saved.

Apart from causing large amounts of human suffering, and epidemic can also destabilize a country and cause large amounts of economic damage (Bloom, Cadarette, & Sevilla, 2018). It is therefore desirable to stop an epidemic as soon as possible. Therefore, minimizing the amount of time before an epidemic is brought under control is also an objective.

Since the WHO is dependent on humanitarian organisations to deliver (medical) aid on the ground (Harman & Wenham, 2018), the WHO objectives must take into account the humanitarian principles on which these organisations operate.

The fundamental humanitarian principles are humanity, impartiality, neutrality and independence (Hilhorst, 2005). Given the scope of the proposed model, the principle of impartiality is the most important to consider. This principle dictates that the provision of aid should be based only on the need that people have. Therefore, people with similar needs should receive the same aid regardless of their background or location. In terms of model objectives, this can be defined as the objective of maximizing equity. In resource allocation equity can be defined in terms of difference between the ratio of met and unmet demand or in terms of difference between demand onset and arrival times of aid (Huang, Jiang, Yuan, & Zhao, 2015). In this research, both definitions will be used. Minimizing the difference between the ratio of met and unmet demand is seen as the most direct translation of the impartiality principle. But large differences in arrival times can cause great human suffering, and should therefore also be considered.

Summing up, the objectives that will be considered are:

- Effectiveness - Operationalised as the number of deaths prevented compared to a base scenario, with no response.
- Efficiency - Operationalised as minimizing the cost in dollars per death prevented.
- Speed of the response - Operationalised as minimizing the number of timesteps until 70% of the infected population is receiving treatment.
- Equity in met demand - Operationalised as minimizing the difference between the ratio of met demand (patients admitted to an ETC) and unmet demand (infected persons not admitted) between regions.
- Equity in arrival times - Operationalised as minimizing the difference between onset of demand (the first occurrence of an infection) and the arrival of aid (an ETC becoming operational) between regions.

Effectiveness and efficiency (or closely related formulations) are common objectives on which resource allocation strategies for epidemics are optimized (see Chapter 1) - the other three are introduced here in order to represent the humanitarian perspective.

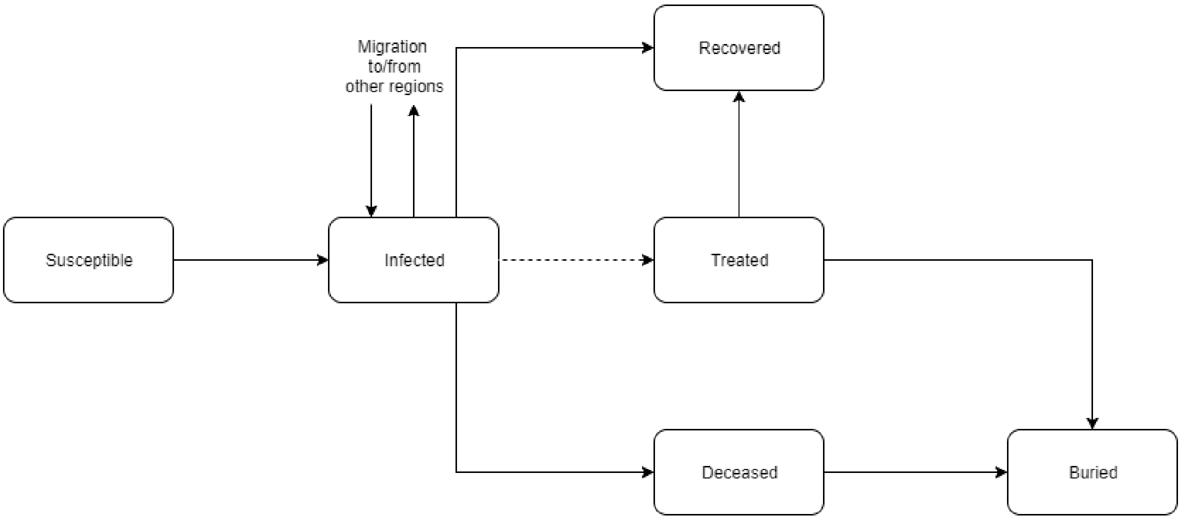

Figure 3.2: Compartmental Model Structure (adapted from Büyüktahtakın et al. (2018))

3.3. Basic Model Structure

3.3.1. Compartmental Model

Compartmental models are widely used in epidemiology to study the dynamics of a disease. They consist of several compartments (i.e. Susceptible, Infected, Recovered) and rates of flow between compartments that are determined by the levels of the compartments themselves (which means it essentially represents a system of differential equations) and epidemiological parameter such as the transmission rate and the infectious period (Brauer, 2008).

Since epidemiological modelling is a science in itself, this research will draw on existing compartmental models for Ebola. As a basis, the model by Büyüktahtakın, des-Bordes, and Kıbış will be used. They developed a deterministic Ebola model to study the impact of intervention starting dates and the optimal spatial placement of ETCs (Büyüktahtakın et al., 2018). The model consists of six compartments: Susceptible, Infected, Treated, Recovered, Deceased, and Buried (where Deceased represents a compartment for individuals who have died but are not safely buried, and therefore serve as a source of infection). Figure 3.2 shows the compartmental model and flows between compartments. In the model by Büyüktahtakın et al. (2018) those who pass away while in the treated compartment first flow into the deceased compartment before being safely buried. However, it is assumed that responders operating ETCs highly prioritize safe burial as a way to stop the disease, and therefore all patients who die at an ETC are immediately safely buried.

3.3.2. Decision Levers

Each time step (i.e. every week), a decision can be made to place one ETC(s) in a district. An ETC can have a capacity of 10, 50 or a 100 beds, and will become operational after a number of time steps depending on its size. Once the ETC is operational, the capacity of the Treated compartment is increased with the number of beds of the ETC. Each time step, it is also possible to hand over one or more existing ETCs to local staff if there are no more patients in the area.

Instead of placing an ETC, the decision-maker can also choose to send one or more surveillance teams to regions. If Ebola is known or suspected in an area, guidelines dictate that surveillance teams are set up to carry out active case finding (World Health Organization and others, 2014). The only function of these surveillance teams is to reduce uncertainty on the state of the epidemic in a certain region.

Uncertain Factor	Description or Range
Number of Patients	• MSF reported hiding of cases, and also stated *“The Ministry of Health and the partners of Kenema hospital refused to share data or lists of contacts with us, so we were working in the dark while cases just kept coming in”* (Médecins Sans Frontières, 2015). • Dutch NGOs reported *“high levels of distrust of international organisations in many communities that led to risky behaviours such as hiding ill relatives, and a reluctance to use health facilities.”* (M’Cormack-Hale, Lavali, & Magbity, 2016). • Underreporting estimated at 250% by a study from August 2014 (Meltzer et al., 2014). • Underreporting estimated at approximately 17%, 70% maximum by a study from December 2014 (Scarpino et al., 2014).
Transmission Rate	• The WHO estimated basic reproduction numbers for the three countries in October 2014: 95% CI of 1.44 to 2.01 for Guinea, 95% CI of 1.72 to 1.94 for Liberia, and 95% CI of 1.79 to 2.26 for Sierra Leone (WHO Ebola Response Team, 2014). • A study from December 2014 gave a basic reproduction number of 1.29 (with a 95% CI of 1.27–1.37) (Scarpino et al., 2014). • A 2015 study on the effect of control measures using fitted models estimated median transmission rates varying from 0.06 to 0.6 depending on the district, with rates as high as 1.5 within the 95% CI (Kucharski et al., 2015). • Superspreading events are also considered as an uncertain factor related to transmission rates. Examples are a burial ritual in Kenema, Sierra Leone leading to 385 secondary causes (Vetter et al., 2016), and in another instance a single case resulting in 24 secondary cases (WHO Ebola Response Team, 2016).
Travelling Behaviour	• A 2016 study by the WHO states that people infected with Ebola can travel quickly and through multiple countries. They highlight an example of a patient fleeing from an ETC and causing secondary cases over 200km away. *“...the sites where new cases would be found, or where transmission would persist, were largely unpredictable”* (WHO Ebola Response Team, 2016). • MSF stated that *“This time, people moved around much more and Ebola travelled with them”* (Médecins Sans Frontières, 2015). • A study from 2016 estimates that 4% to 10% of infected cases migrate to another district, and that of this group, between 0% and 23% migrate to another country (Backer & Wallinga, 2016).
Effect of an ETC	• Fatality rates at ETCs varied significantly: MSF has reported ranges between 36% and 60% (McNiel, 2015). • MSF has also reported that *“In Ebola outbreaks, health facilities without proper infection control often act as multiplying chambers for the virus”* (Médecins Sans Frontières, 2015).

Table 3.1: Table outlining the four uncertain factors.

3.4. Uncertain Factors

The role of uncertainty in humanitarian logistics, and uncertain factors that are studied in the literature of this field, are discussed by Liberatore et al. (Liberatore et al., 2013). They outline several uncertain parameter categories: demand (in terms of size and type), availability of supplies, affected areas, demand location, and the state of the transportation network.

Only uncertain parameters directly related to the facility location decisions are discussed here. Therefore, uncertainties relating to actor roles, authorities and coordination structures, which also played an important part in the response (see for example Karlsen and Kruke (2018)), are not discussed here.

Given the model structure as described in Section 3.3, the following categories of uncertainty are considered to be relevant: demand size, affected areas, and demand location. Demand type is assumed to be known as best practices for treating Ebola were already established. The state of the transportation network is an important and uncertain factor ((Médecins Sans Frontières, 2015) (Interview A) but out of scope for the model.

To identify the uncertain factors relevant for the case study, a variety of sources were used: reports written or commissioned by organisations active in the response (such as MSF, Save the Children, and the WHO), coordination documents by the WHO and UNMEER, newspaper articles, and scientific papers published during or after the epidemic. Four key uncertain factors were identified: uncertainty in the number of patients, uncertainty on the transmission rates, uncertainty in travelling behaviour, and uncertainty on the effect of an ETC on key epidemiological parameters. These factors, along with descriptions of how the uncertainty manifests itself (including variable ranges) are shown in Table 3.1

The effect of an ETC on the disease progression in a region is hard to define on a detailed level based on the available information, and sufficient quantitative data was not available either. Therefore, this factor is not included in the model for the remainder of the research.

3.5. Conclusion

By determining the problem owner and the basic model structure, the model objectives and uncertain factors could be identified. The answers to the sub-questions considered in this chapter are therefore as follows:

1. What are the objectives on which decisions for treatment centre placement are optimized?

The identified objectives are Effectiveness, Efficiency, Speed of the response, Equity in met demand, and Equity in arrival times. These objectives ensure that a variety of perspectives are represented in the model: Affected populations benefit from an effective and speedy response, whereas a responding organisations also consider costs as an important (limiting) variable. Humanitarian concerns are addressed with the two equity objectives.

The second sub-question

2. What are the main sources of uncertainty relevant fo a treatment centre placement decision?

Was answered on the basis of academic literature, reports by responding organizations, newspaper articles and expert interviews. These resulted in the identification of the following uncertain factors relevant for the placement of an ETC: The number of infected individuals in a region, the transmission rate, the rate of travel between regions, and the effect an ETC has on the disease progression in a region. Because the last uncertain factor could not be defined or quantified to the degree necessary for a simulation model, this factor is not incorporated in the remainder of the research.

4 Uncertainty Reduction

An important goal of this research is to investigate how reduction of uncertainty through decisions influences optimal policies for treatment centre placement. This chapter aims to formulate an answer to the third sub-question:

> 3. How can the reduction in uncertainty resulting from a treatment centre placement decision be conceptualized?

A core assumption is that making a decision to place resources in a region where the intensity of the epidemic is uncertain, can be valuable later in the response. A presence in the region will mean that this uncertainty is reduced, and may lead to better decision-making in the long term. Uncertainty is reduced as information is gathered from patients in the ETC, contact tracing takes place, and the number of patients itself is also an indication of how the epidemic is evolving.

This chapter provides a conceptualisation of this uncertainty reduction in Section 4.1, which also explores which factors contribute to the reduction of the uncertain variables identified in Chapter 3.4. In order to develop policies that make strategic use of uncertainty reduction, is is necessary to distinguish between decisions aimed primarily at uncertainty reduction and decisions aimed at exploiting current knowledge. This is done in Section 4.3, which classifies explorative versus exploitative decisions.

4.1. Conceptualisation of Uncertainty Reduction

The increase in situational awareness, and therefore the decrease of uncertainty, can be represented as a function over time. This function then represents the percentage with which the range of uncertainty is reduced over time. This way of representing an increase in knowledge is based on information theory, in which functional forms are used to represent how information diffuses through a network (see for example Yang and Leskovec (2010)).

This approach requires the functional form of the uncertainty reduction to be determined. However, no quantitative data on the level of uncertainty and its reduction over time on variables such as transmission rates or the number of infected people in an epidemic exists. As a result, the functional form cannot be derived quantitatively and instead will be approximated using qualitative data. Additionally, the functional form of the uncertainty reduction may differ depending on which variable it represents. Therefore, for each of the uncertain variables, the relevant factors on which the functional form depends will be discussed below. Based on this analysis the functional form will be determined for each of the variables.

The reduction of uncertainty as discussed in the next three sections is based on the decision to place an ETC in a region. The effect of sending only a surveillance team to a region is discussed separately afterwards.

4.1.1. Number of Infected Individuals

Relevant Information Sources

Uncertainty on the number of infected individuals is reduced by the following information sources:

- **Number of patients in an ETC**: The number of patients at an ETC is assumed to be proportional to the total number of infected individuals. If an ETC had no free beds, people were known to queue outside (Nyenswah et al., 2014).

- **Information gained from contact tracing**: As part of the standard Ebola response to a suspected, probable or confirmed case being identified (which may be done by mobile teams or at an ETC), contact tracing will take place. In this process, all persons who may have been infected by the case are followed for 21 days to see if they develop symptoms (World Health Organization & Centers for Disease Control and Prevention, 2015). If they do not develop symptoms within this timeframe, it is concluded they do not have Ebola. The incubation period for Ebola is estimated at 11.4 days (WHO Ebola Response Team, 2014). Therefore, contact tracing results in three moments in time in which uncertainty is decreased: within days, the number of contacts (and therefore the maximum number of secondary cases resulting from the patient) becomes known. After 1 tot 2 weeks, secondary cases will start to present and are immediately known. Finally, after 3 weeks it can be said with certainty that those who have not developed symptoms do not have Ebola.

- **Information gained from (informal) contact with locals**: The assumption is that over time, responders will gain information about the local situation by talking to locals and gaining their trust. This information could come from patients inside the ETC and their family or from local staff or community leaders.

Functional Form

A first assumption on the function representing the reduction of uncertainty on the number of infected persons in a region is that uncertainty cannot be reduced to 0%. Hiding of cases due to distrust was a phenomenon reported by several NGOs (M'Cormack-Hale et al., 2016; Médecins Sans Frontières, 2015), and estimates on the number of unreported cases range from 17% to 250% (Meltzer et al., 2014; Scarpino et al., 2014). Case tracing was sometimes impossible due to the high number of cases and transmission chains (Adams et al., 2016).

The next assumption is that the factors described above do not reduce uncertainty more if more patients are seen. This is due to the fact that the number of infected individuals varies over time, and therefore new information needs to be gathered continuously. It is useful to think of each of these factors as an information channel that becomes available over the period of time an ETC is set up and becomes operational.

The following assumptions are therefore made to determine a functional form. First, that once a decision is made to open an ETC and construction begins, information will be gained from contact with locals. Once the ETC becomes operational, the number of patients admitted (or turned away) is known. Over the course of the next three weeks, information from contact tracing of the first patients comes in, and this process carries on continuously as new patients are admitted. Over the course of these three weeks, it is also assumed that the trust of the local population is won and they share whatever information they know. As a result, three weeks after an ETC becomes operational, all available information channels are exploited maximally.

The function representing the reduction of uncertainty on the number of infected people is therefore only dependent on time. The proposed functional form is shown in Figure 4.1: A sigmoidal function represents the uncertainty reduction as a result of opening an ETC.

The sigmoidal shape chosen for the opening of an ETC is argued for as follows: while an ETC is being built, local workers as well as curious visitors provide some information on the current status in the region, based on their own experiences. Right after the ETC becomes operational, uncertainty is reduced at the highest rate as patients come in and contact tracing is set up. After three weeks the first cycle of contact tracing is complete and most information channels are now exploited fully. As the local

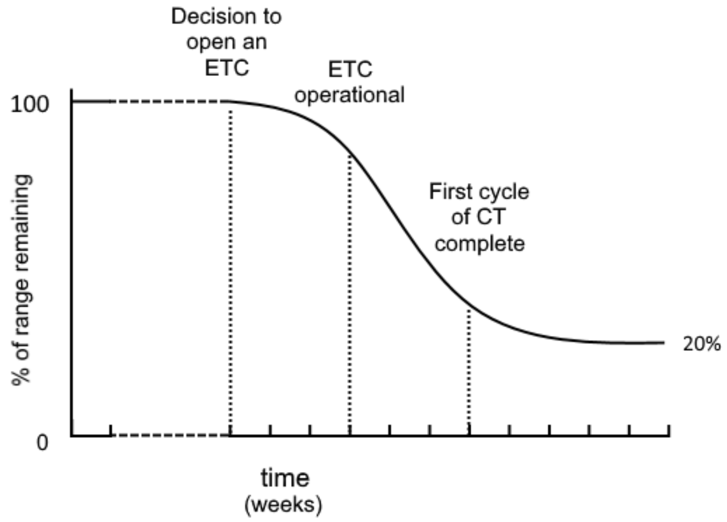

Figure 4.1: Functional form for the reduction of uncertainty on the number of infected individuals in a region as the result of a 50-bed ETC being placed in that region.

population gains more trust in the responders, they may start providing some additional information in the next few weeks.

4.1.2. Transmission Rate

Uncertainty on the transmission rate in a region is reduced by the following sources:

- **Rate at which patients are admitted at an ETC**: As for the number of infected indivuals, it is assumed that the number of patients at an ETC is proportional to the total number of infected individuals in that region. The rate at which patients are admitted (or turned down) is therefore indicative of the rate the disease is spreading.
- **Information gained from contact tracing**: In its basis, the transmission rate depends on the probability of infection upon contact, and the rate of contact between infectious individuals and the susceptible population. With contact tracing, the number of people an infectious individual has had contact with becomes known within days. Assuming the total population of a region is known with reasonable accuracy, the rate of contact can then be estimated with increasing accuracy for every patient for which contact tracing is carried out.

Functional Form

For both relevant sources, it can be argued that as the number of observations increase, the transmission rate can be estimated more accurately. If it is then also assumed that the pre-admittance behaviour of individuals who seek treatment is representative over the whole infected population, uncertainty on the transmission rate could approach 0 over time.

It is therefore proposed that, when an ETC becomes operational, the reduction of uncertainty occurs as an exponential function of the number of patients as shown in Figure 4.2. An exponential function is chosen as at the beginning, gaining factual information about patient behaviour helps to obtain much better estimates for the transmission rate, and as such uncertainty is quickly reduced. As the cumulative number of patients increases, the impact one patient has on improving the estimate becomes smaller and the curve flattens out.

4.1.3. Travelling Behaviour

Uncertainty on the travelling behaviour is reduced by the following sources:

- **Information gained from contact tracing**: In the process of contact tracing, a patient will be asked about their contacts and movements. In this way, it will become apparent if a patient has travelled from another district while they were infectious.

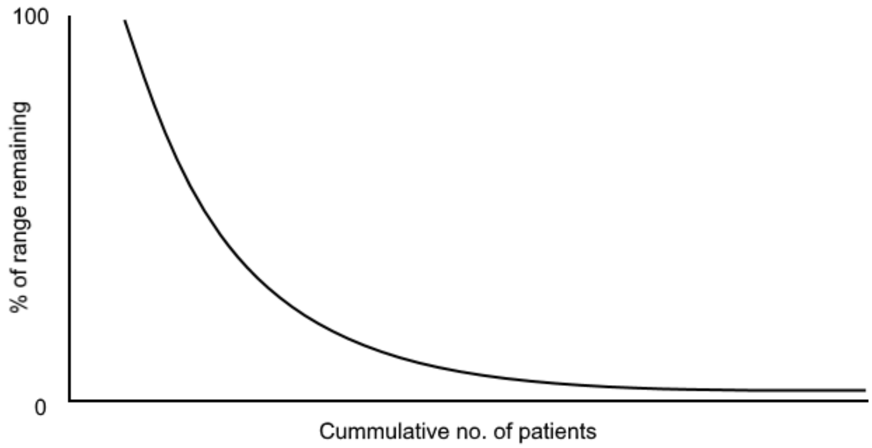

Figure 4.2: Functional form for the reduction of uncertainty on the transmission rate.

- **Information gained from (informal) contact with locals**: As trust between responders and locals and patients grows, they may share information about their regular travelling patterns to other districts, for example for trading, work, or to visit family.

Functional Form
As for the transmission rate, it is assumed that the behaviour of admitted patients is representative of the whole infected population. Therefore, travelling behaviour can be estimated more accurately as the number of patients seen increases. Information from locals could theoretically become available as soon as construction of an ETC begins, and is only dependent on time. However, it is assumed that trust is not strong enough yet to gain this information right before or after an ETC is opened. It is also assumed that during the time it takes to build up this necessary trust, the information gained from informal contact does not significantly add to the obtained information from contact tracing. Therefore, it is proposed that the reduction of uncertainty on travelling behaviour is an exponential function of the number of patients. The functional form is therefore the same as for the transmission rate.

4.1.4. Surveillance Teams

As established in 3.3 a decision maker can also choose to send a surveillance team to a region. Since these teams are mobile and do not establish a permanent presence in the region, many of the information channels discussed above are not available to them. Therefore, it is assumed that the main purpose of surveillance teams is to give a first report of the situation in the region. This is conceptualized by assigning all regions the status of "hidden" at the start of the response, meaning the decision-maker has no information at all (not even an uncertainty range) on the state of the epidemic. Sending a surveillance team to a region can then remove this "hidden" state at which point information about the region becomes available with full uncertainty (i.e. the range is at 100%). An ETC can also be placed in a region to remove the hidden state, with the difference being that it will then continue to reduce the range of uncertainty. Surveillance teams have no effect if they are sent to or active in a region which is no longer hidden.

4.2. Information Delay

The sections above discussed uncertainty reduction from the perspective of someone at a single ETC. However, as established in Chapter 3, the problem owner who is interested in this uncertainty reduction is coordinating the response at an international level. This means that the information needs to be communicated to them from the ETC, which will introduce some time delay. The length of this delay may vary depending on the source of information: one expert (Interview B) noted that for surveillance and contact tracing, it may take days for word of a fatality to reach a district health manager, who then sends this information to the national coordinator via mail, introducing another few days of delay. Another expert outlined that while reports from national coordination organisations would come in daily, they were not standardized, so extracting and preparing this data for international coordination provides another hurdle (Interview D). Delays resulting from reformatting and parsing data from a variety of different formats were also noted by another expert (Interview F). The WHO reported an average delay of 6 days between symptom onset and their notification (WHO Ebola Response Team, 2014). Therefore, an information delay of at least a few days to a week should be introduced. This delay can be represented by shifting the proposed uncertainty reduction functions to the right, thus delaying the point in time at which uncertainty reduction is perceived.

4.3. Explorative vs. Exploitative Decisions

Given the assumption that it can be strategically beneficial to make a decision that reduces uncertainty, there is a need to define what qualifies as an uncertainty-reducing decision and what does not. Borrowing terminology from Reinforcement Learning, uncertainty-reducing decisions will be called explorative decisions, whereas decisions aimed at delivering aid to those most in need will be referred to as exploitative decisions.

4.3.1. Defining explorative and exploitative decision

Since decisions are defined as explorative or exploitative from the perspective of the decision-maker, this perspective first needs to be established. The information the decision-maker has available to them at any time are:

- The values of the known variables in the model (i.e. time from infection to recovery or death).
- The total population of each region.
- Ranges (i.e. min. and max. values) on the uncertain factors for each region where Ebola is known.
- The number of patients in an ETC in a region (with some time delay).
- Knowledge on how uncertainty will be reduced as the result of a placement decisions (i.e. the functional forms discussed in the previous section).

of which the uncertain factors for the number of infected individuals and the transmission rate are the most important indicators of the severity of the epidemic. In interview A, the interviewee indicated that there would be attempts to forecast demand in order to inform placement decisions, but that this was very difficult, which supports the assumptions made here.

It is useful to define explorative and exploitative decisions in relative terms rather than with absolute ones. Absolute terms, such as defining a decision as explorative when the uncertainty in that region has not been reduced by more than $x\%$, can lead to a standstill when such a definition applies to all or none of the regions. This could result in the distinction becoming useless at the very start and at the later stages of a simulation run.

Instead, if the purpose of the distinction between exploration and exploitation is kept in mind, and the regions are considered relative to each other, the definition becomes straightforward. An explorative decision is made to learn more about regions of which very little is known. Therefore, if we want to make an explorative decision, we should choose the region with the highest level of uncertainty. If several regions have the same level of uncertainty, one can be picked at random. Similarly, an exploitative decision is made to serve the humanitarian imperative: to provide aid to those most in need. Therefore, an exploitative decision takes place in the region with the highest guaranteed number of beneficiaries (i.e. the region with the highest lower bound on the predicted number of infected individuals).

Since the decision maker also knows the range in which the transmission rate lies, the evaluation for an exploitative decision could be enhanced by considering the estimated number of infections at the timestep in which an ETC would become operational, instead of the current number of infections.

4.3.2. Choosing the type of decision

In the preceding section, explorative and exploitative decisions are defined by prescribing what region to choose given the type of decision we want to make. So how is the type of decision to be taken in a particular timestep determined? This will be dictated by a policy function dependent on the total level of uncertainty in the system. The shape of the policy function is given at the start of the simulation run and is based on the results of the direct policy search. It returns the probability p with which an explorative action should be taken. The type of decision taken that timestep is then an explorative decision with a probability of p and exploitative with a probability of $1 - p$.

4.4. Conclusion

This chapter has served to answer the third sub-question:

> 3. How can the reduction in uncertainty resulting from a treatment centre placement decision be represented conceptualized?

The reduction in uncertainty can be represented in terms of the percentage of the original range around the uncertain variable that remains. By studying the factors that provide information about the uncertain variables a functional form that shows how uncertainty is reduced can be determined. This function can depend only on time or also on factors such as the cumulative number of patients. In order to use the idea of uncertainty reduction strategically a distinction needs to be made between decisions made to reduce uncertainty, which are referred to as explorative decisions, and decisions made to immediately relieve suffering, which are called exploitative decisions. For an explorative decision, a surveillance team or treatment centre is placed in the region with the highest level of uncertainty. For an exploitative decision, the region with the highest number of guaranteed beneficiaries is chosen. The ratio of explorative versus exploitative decisions that performs best is determined by a policy function, which can be found through direct policy search.

5 Model Design

In this chapter, the requirements for the model resulting from the previous chapters, as well as the relevant assumptions made for its design, are outlined. This is done with the purpose of translating the conceptual model developed thus far into a simulation model, thereby answering the fourth sub-question of this thesis:

> 4. How can the sequential treatment centre placement decision problem and the decision-uncertainty reduction be represented in a simulation model?

First, the key components and behaviours identified as relevant for the model in the previous chapters are formulated as model requirements. Because a simulation model can never be as complex in scope or structure as the real system, the most relevant assumptions are listed in the following section. The remainder of the chapter describes how the key model components are implemented.

5.1. Model Requirements

Based on the previous chapters, a series of requirements the simulation should adhere to can be formulated. The most important requirements are summarized in Table 5.1 . They are separated into three categories, one relating to the epidemiological and facility location structure which forms the basis of the model, the second related to decision-making, and the third category concerns uncertainty in the model.

5.2. Assumptions

In order to realize an operational simulation model, some assumptions have to be made - either to simplify reality or in order to represent phenomena which are known to exist, but whose mechanisms are unknown. Key assumptions are:

- Births are not considered in the model. The duration of the epidemic that is modelled is deemed short enough for births to not have a considerable effect on the population. Deaths from other causes than Ebola are also not included.
- Only the travelling of infected individuals is assumed to have relevant effects. Therefore, travel behaviour of individuals in other compartments is not implemented. This means that dead bodies being moved to different regions is also not implemented.
- The travelling behaviour of infected individuals is assumed to be constant - i.e. the rate of travel does not change over time or due to certain effects. This also means that behavioural effects such as migrating away from disease-ridden regions are not considered
- For simplicity, individuals can only travel to regions which are immediate neighbours of their region - unless they travel in the context of a superspreading event in which case they will travel to a non-neighbouring region.

Epidemiological and Facility Allocation Model Structure The simulation model needs to...
• be able to represent the spatial spread of the epidemic over time. • keep track of the number of individuals in a compartment in a region. • incorporate travel of infected individuals to neighbouring regions. • incorporate superspreading events where an infected individual travels to a non-neighbouring region. • translate the impact of the placement of resources (an ETC or surveillance teams) to the model behaviour in a specific region.
Decision-Making The simulation model needs to...
• distinguish between explorative and exploitative decisions. • determine what type of decision to make based on the total level of uncertainty present in the system. • incorporate the decision levers which have been identified in Chapter 2. • be able to do nothing. • be able to remove resources from an area. • keep track of which resources are used, where they are used, and when they are used. • constrain the number of resources available for allocation. • keep track of the data necessary to evaluate the objectives defined in Chapter 2.
Uncertainty The simulation model needs to...
• incorporate the uncertain factors identified in Chapter 2. • keep track of ground truths which dictate how the epidemiological model evolves over time. • maintain a a range of values around the ground truths of regional variables, dependent on the level of uncertainty, which is available for decision-making. • keep track of the overall level of uncertainty in the model as experienced by the decision maker. • keep track of the level of uncertainty experienced by the decision maker in a region. • distinguish between regions where no epidemiological information is available, and regions where epidemiological information is available (with some uncertainty). • follow the functional forms determined in Chapter 3 for the reduction of uncertainty around variables.

Table 5.1: showing the model requirements, separated into three thematic categories.

- Each timestep, a decision can be made to place one ETC, send out a maximum of three surveillance teams, or to do nothing. Allowing only one ETC to be placed per timestep is a simplying assumption, though it is also supported by a statement from an interview subject who commented on the difficulty of finding organisations willing to run an ETC (Interview C), which suggests operationalizing large amounts of capacity at once was difficult. The maximum of three surveillance teams is an arbitrary assumption rationalized by the idea that it will then take 6 weeks to explore a grid of 16 regions, equivalent to a country the size of Sierra Leone.
- Surveillance teams can only reduce uncertainty by revealing the state of a “hidden” region of which previously no information was known at all. Therefore, surveillance teams can only be sent to regions which are hidden, as they would have no effect otherwise.
- It is assumed that an ETC with a certain capacity can be placed anywhere for the same cost.
- It is assumed that an ETC with a capacity of 10 beds takes 1 week to construct and become operational, an ETC with a capacity of 50 beds will take 3, and an ETC with a capacity of 100 beds will take 4 weeks, based on Sánchez Carrera (2015).
- It is assumed that all beds in an ETC are filled with infected individuals. This was not the case during the Ebola epidemic and in fact, beds being occupied by patients who had not received tests results was identified as a bottleneck issue by one of the interview subjects (Interview E). However, the issue of streamlining diagnosis is seen as a separate problem from the resource allocation problem, and is therefore ignored for the sake of simplicity.
- ETCs reduce uncertainty of variables according to the functional forms defined in Chapter 4.
- Due to travelling behaviour, if an ETC becomes operational, uncertainty in neighbouring regions is reduced to 95%.
- The level of uncertainty is determined by whichever functional form relevant to that region provides the most reduction, that is, uncertainty reduction is not additive.
- All infected individuals admit themselves to an ETC if beds are available.
- If there are more infected individuals than free beds in a region, the infected individuals who do not move to the treated compartment flow to the recovered or treated compartment at the same rate as they would if there was no treatment centre at all.
- The time that an individual has spent in the infected compartment does not influence the time the individual spends in the treatment compartment - theoretically an individual could spend two weeks in the infected compartment, and then spend another two and a half weeks in the treated compartment before they have recovered. This is because a system dynamics model does not keep track of individual cases.
- All individuals who die in an ETC are safely buried. This is not the case in the model by Büyüktahtakın et al. (2018). However, based on reports and literature by medical humanitarians, safe burial was identified as having high priority for responders. Therefore, it is assumed that the responders running an ETC will ensure safe burial of all deceased patients.
- ETCs can be handed over to local staff. The ETC must have been open for at least two weeks, as it is assumed local staff need to be recruited and trained. Once the decision is made to close an ETC, it takes 1 week for a small ETC (10 beds) to have handed over operations and resources are freed again, and 2 weeks for a big ETC (50 or 100 beds).
- The transmission rate in a region is constant over time. This simplifying assumption is unrealistic as it does not account for changing behaviour resulting from fear or community outreach programmes. However, it is believed that the epidemiological dynamics relevant to the problem formulation of this thesis can still be modelled to an acceptable degree with a constant transmission rate.

5.3. General Structure

The simulation model has three key components that work together: a compartmental model that simulates the progression and geographic spread of the disease over 16 regions which is shown in Figure 5.1, a set of objects and functions that keep track of the uncertainty related to regional variables and update them as they are reduced, and a decision-making module that decides where to place resources given the current state of the system. Each of these three components will be discussed in more detail below.

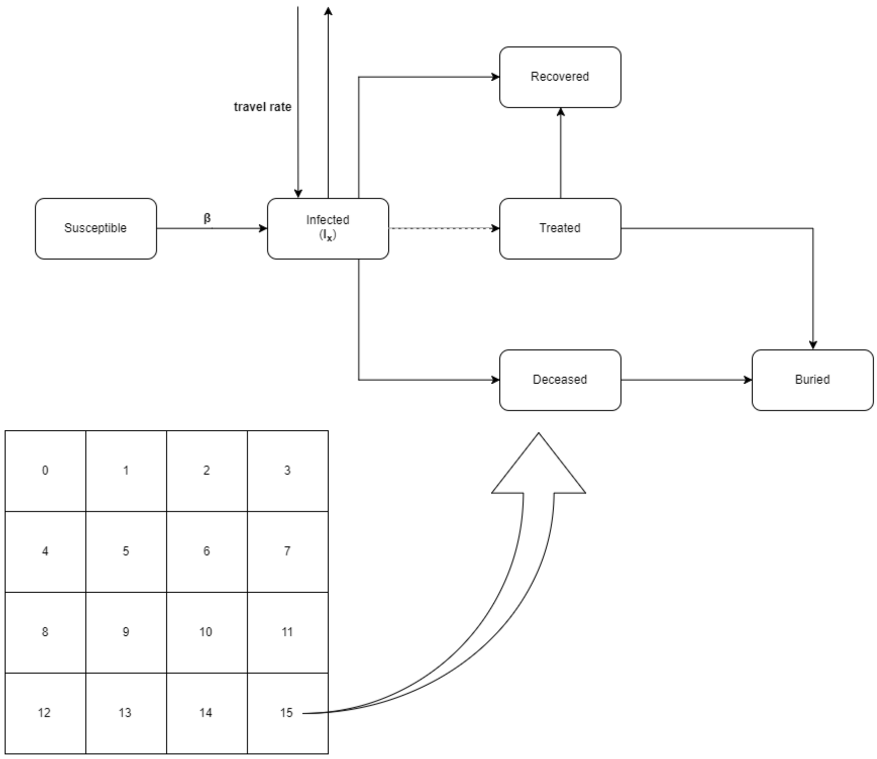

Figure 5.1: In each of the 16 regions, a compartmental model keeps track of the progression of the epidemic. In this way, the spatio-temporal development of the epidemic is modelled.

5.3.1. Compartmental Model

The regions are represented with a square grid, in which every cell represents a region. The progression of the disease in a region is simulated using a compartmental model, as discussed in Section 3.3.

Individuals move from the Susceptible (S) compartment to the Infected (I) compartment with the rate:

$$\frac{dS}{dt} = -(\beta_i * I + \beta_d * D) * \frac{S}{S + I} \tag{5.1}$$

where β_i represents the transmission rate of infected individuals, and β_d the transmission rate of deceased individuals and D represents the number of individuals in the Deceased compartment, who have not been safely buried and are therefore still infectious.

Infected individuals can move to three different compartments: the Deceased compartment if they die, the Recovered compartment (R) if they survive, and if there is free treatment capacity they move to the Treated (Tr) compartment. Additionally, infected individuals can leave the region by travelling, and infected individuals from neighbouring regions can travel into the region. Therefore, the rate of change for the infected compartment is given by:

$$\begin{aligned}\frac{dI}{dt} =& (\beta_i * I + \beta_d * D) * \frac{S}{S+I} \\ &+ infected\ travellers \\ &- r_{travel} * I \\ &- \frac{(1 - f_{\text{without treatment}}) * I}{t_{\text{recovery without treatment}}} \\ &- \frac{f_{\text{without treatment}} * I}{t_{\text{death without treatment}}} \\ &- min(ETC\ capacity - Tr, I)\end{aligned} \tag{5.2}$$

where r_{travel} is the rate of travel, $f_{\text{without treatment}}$ is the fatality ratio without treatment, $t_{\text{recovery without treatment}}$ is the time to recovery without treatment in weeks, and $t_{\text{death without treatment}}$ is the time to death without treatment in weeks.

Infected individuals move immediately to the treatment compartment if there is free capacity. From there, they can either move to the recovered compartment or to the Buried (B) compartment, meaning that if they do not survive they are immediately safely buried:

$$\begin{aligned}\frac{dTr}{dt} =& min(ETC\ capacity - Tr, I) \\ &- \frac{(1 - f_{\text{with treatment}}) * Tr}{t_{\text{recovery with treatment}}} \\ &- \frac{f_{\text{with treatment}} * Tr}{t_{\text{death with treatment}}}\end{aligned} \tag{5.3}$$

Given the inflow from the Infected and Treated compartments, the rate of change for the Recovered compartment is as follows:

$$\frac{dR}{dt} = \frac{(1 - f_{\text{without treatment}}) * I}{t_{\text{recovery without treatment}}} + \frac{(1 - f_{\text{with treatment}}) * T}{t_{\text{recovery with treatment}}} \tag{5.4}$$

Since those who die in the Treated compartment are immediately (safely) buried, the rate of change in the Deceased (D) compartment is only dependent on the number of infected people and the rate at which safe burials are conducted:

$$\frac{dD}{dt} = \frac{f_{\text{without treatment}} * I}{t_{\text{death without treatment}}} - D * r_{\text{safe burial}} \tag{5.5}$$

Finally, the Buried compartment then receives inflows from the Deceased and Treated compartments:

$$\frac{dB}{dt} = D * r_{\text{safe burial}} + \frac{f_{\text{with treatment}} * Tr}{t_{\text{death with treatment}}} \tag{5.6}$$

Random Travelling

As described in Chapter 2 and in the requirements, super-spreading events can occur in the form of an infected individual travelling to a non-neighbouring region. This can lead to the epidemic spreading to regions where it was previously not present.

These super-spreading events as a result of random travelling were implemented in the following way: it is assumed that approximately 1 in 200 infected individuals travels long distance, based on Backer and Wallinga (2016). As a logical consequence, the number of super-spreading events is dependent on the total number of infected individuals in the system N_I. Each timestep, the number of random travellers is determined using the following function:

$$Travel_{random} = N_I * 0.005 \tag{5.7}$$

This formula does not provide an integer number of travellers. Therefore, the number of travellers is determined by rounding the result down, with a probability of one additional traveller equal to the fractional digits of the result [1]. For each of the travellers, the region of origin is determined at random under the condition that the region has at least one infected individual in it. The destination is also chosen at random, under the condition that it is not the region of origin or one of its neighbouring regions.

Since the removal of single individuals from regions and placing them in another is a discrete event, this is implemented with a function after the compartmental model is run. The list that keeps track of the states of all the regions is then updated accordingly and passed to the compartmental model in the next iteration (timestep).

5.3.2. Implementation of Uncertainty

Uncertainty for the Decision Maker

At the start of a simulation run, all regions are “hidden”, meaning that there is no epidemiological information available from the regions for decision-making. There are two ways in which information from a region can be made available. The first is to make an explorative decision to send in a surveillance team or to place an ETC. If a surveillance team is sent in, epidemiological information becomes available but with the maximum range of uncertainty. For an ETC, epidemiological information becomes available and the uncertainty is reduced over time according to the functional form as described in Chapter 3. The exact values used in the simulation model are outlined in Appendix B. The second way in which a region can become “unhidden” is by chance. With some probability, depending on the number of infected individuals, the epidemiological information will become available with the maximum range of uncertainty. This represents the idea that the more severe the situation in a region is, the higher the chance is that news about this travels to the decision-maker level. This probability is given by:

$$p = 1 - \frac{1}{1 + exp(\frac{I-n}{z})} \tag{5.8}$$

and is referred to as the “spontaneous news” mechanism in the remainder of this thesis. The function has a sigmoidal shape, and approaches the limit of 1 at around $2n$ infected individuals I. The variable z determines the steepness of the slope. For the simulation model, $n = 40$ meaning that at 40 infected individuals the chances of receiving spontaneous news from that region are 0.5. $z = 7$, reflecting that at around 20 infected individuals, the chances of receiving news start to increase more steeply. These values are based on assumption and aim to hit a balance between densely populated urban areas from which news travels much faster, and rural areas which are more isolated. Since no distinction between urban and rural areas is made in the model, this function is used in all regions.

As stated in the requirements, the ground truth of uncertain variables is not available for decision making. Instead, a range around the ground truth is known, where the width of this range is dependent on the current level of uncertainty for that variable in a certain region. This was implemented by creating objects for uncertain variables, which keep track of both the ground truth and the level of uncertainty. If the variable is accessed for decision making, the appropriate range around the ground truth is given. The distance between the ground truth and the bounds of the range is random, that is, no assumption is made that the knowledge of the decision maker symmetrically converges to the ground truth. Instead, each time uncertainty is decreased for a variable, the respective range is initialised randomly around the ground truth.

[1] For example, if there are 243 infected individuals in the system, the result of the function above is 1.215. There is then at least 1 traveller, and with a probability of 0.215 there are 2.

5.3.3. Decision-making Module

The decision-making module calls either the function to make an explorative decisions or the function to make an exploitative decision. Which function is called is determined by the policy function.

Policy Function

The policy function returns the probability with which an exploitative action should be taken based on the current total level of uncertainty in the system. The function is represented by two cubic radial basis functions following Quinn et al. (2017). The probability of taking an exploitative action given the total level of uncertainty U_i in the system is:

$$p_i = min(max(w_1 * |\frac{U_i - c_1}{r_1}^3| + (1 - w_1) * |\frac{U_i - c_2}{r_2}^3|, 0), 1) \tag{5.9}$$

The type of decision taken that timestep is then an explorative decision with a probability of p_i and exploitative with a probability of $1 - p_i$.

The shape of the policy function is determined by the radial basis centres c_1, c_2, its radii r_1, r_2 and the weights given to each of the basis functions, determined by w_1. These variables are defined with the model initialization and allow for different policies to be run.

Explorative Decisions

For an explorative decision, the algorithm first finds the regions with the highest level of uncertainty. The level of uncertainty of a region is determined by adding up the percentages of uncertainty around each of the uncertain variables (infected individuals and the transmission rate). If a region is hidden, these percentages are assumed to be 150%. If there are multiple regions with the same level of uncertainty, one or more (depending on the number of free resources) are chosen at random. If the region of choise is still "hidden", the first choice is to send a surveillance team. If there are no free surveillance teams or if the region is not hidden, a small ETC is placed. If no resources are free, nothing is done.

Exploitative Decisions

For an exploitative decision, the algorithm finds the region with the highest estimated number of patients for the next timestep. This is a risk-adverse estimation based only on the total population of a region and the lower bound on the current range for infected individuals. The choice was made to work with this risk-adverse estimation since humanitarian organisations and the WHO need to justify the spending of resources to their donors. Using the upper bound for the number of infected individuals might result in more resources being sent than there is need for, whereas the risk adverse approach guarantees there is at least sufficient need, if not more. If multiple regions have the same number of predicted infections, one is chosen at random. If the projected number is higher than 50, an ETC with a capacity of 50 is placed, and for 100 or more projected cases one with a capacity of 100. If the resources for this are not available, or if the projected number is lower, a small ETC is placed (with a capacity of 10). If no resources are available, or if the decision-maker is not aware of any cases, nothing is done.

5.3.4. Objectives

The exact implementation of the model objectives is provided in Appendix C. Here it is noted that in the simulation model, and the subsequent results discussed in the following model, when referring to specific values obtained for some of the objectives, the objectives are formulated differently. This is done to avoid confusion when discussing model results, because these objectives should be minimized:

Efficiency is referred to as **Cost per Death Prevented**

Equity in Met Demand is referred to as **Difference in Met Demand**.

Equity in Arrival Time is referred to as **Difference in Arrival Time**.

Speed is referred to as **Time until Containment**.

The name of the **Effectiveness** objective remains the same.

5.3.5. Implementation

The model was implemented using Python 3. The compartmental model was implemented as a set of differential equations using ordinary differential equation solver `odeint` from the SciPy library. For most other mechanisms discussed above objects or functions were created to provide the necessary functionality. Appendix I provides an oversight of where in the code specific functionalities can be found. All code can be found on `https://github.com/edenbrok/thesis`.

6
Model Validation

This chapter discusses the validity of the simulation model, as outlined in the previous chapter. It does so from two perspectives: first, the validity of the behaviour of the epidemiological and outcomes of the model is discussed. This is to ensure the simulation model provides a correct representation of the epidemic and the performance of a policy is represented correctly. The next section investigates whether the designed model is valid given the research purpose of this thesis. The chapter closes with a short discussion on the validity of the research approach itself.

6.1. Validity of the Simulation Model

In order for the results to have relevance for an actual response, it needs to be established whether the model parametrization and relevant assumptions lead to realistic model behaviour and outcomes. This will discussed here on two levels of model behaviour: on the level of the epidemiological model, and on the level of the objective scores.

6.1.1. Epidemiological Model

The parametrization of the epidemiological model was taken from Büyüktahtakın et al. (2018), who validated their epidemiological model using case-data from the 2014 Ebola response. However, their model is different from the one in this thesis as they used one compartmental model to represent the entire country of Sierra Leone, while here it is used to represent a district. Using district level case data World Health Organization (2016), the case growth in the absence of a response was checked for 6 timesteps (equivalent to 6 weeks) and found to be accurate. However, when the model is run for 26 timesteps (representing the full period of interest from June 2014 to December 2014) without any response measures implemented, it produces only a fourth of the total cases seen in the actual epidemic (which includes response measures). The reason why Büyüktahtakın et al. (2018) do realize the right number of cases in their simulation model is because their epidemiological model represents all of Sierra Leone, resulting in an higher absolute case growth as a result of the exponential increase in cases.

It is however, not the mechanism in which cases are generated in reality, leaving the question of why the regionalized model does not produce enough cases even though it follows the real-life data for the first few weeks. This is explained when considering super-spreading events. These are included in the model, but only in the dimension of spatial movement. However, there are many reports of one individual causing tens of new cases (Vetter et al., 2016; WHO Ebola Response Team, 2016). This type of superspreading is not implemented in the model, though it would likely lead to more realistic case numbers. The motivation to not include it is that it would be a stochastic mechanism which has a huge impact on all the objectives. That is, it would introduce high-impact stochasticity in the model, meaning that the number of model runs necessary to detect influence from other factors was expected to be unmanageably high. Since all outcome objectives related to the number of cases are expressed in relative terms (i.e. percentage of deaths prevented compared to no response, the difference in cumulative patients over total infections), the difference in simulated and real cases is accepted.

6.1.2. Model Outcomes

The realism of the model outcomes will not be discussed in much detail here, given the attention that these will receive in Chapters 8 and 9. However, one important assumption is outlined here. This concerns the fact that the two equity objectives are calculated on relative terms. That results in the fact that if one region has 10% unmet demand, equivalent to 50 cases, and another has 25% met demand equivalent to only one case, this difference is “punished” equally by the objective as it would have if the 25% had been equivalent to 75 cases. The same holds for equity in arrival times. This is a value-based choice - even though many humanitarian logistics employ more utilitarian methods which would make a distinction between these two cases (Huang et al., 2015). The motivation for this formulation of the equity objectives is based on the humanitarian charter, which states *“...that all people affected by disaster or conflict have a right to receive protection and assistance to ensure the basic conditions for life with dignity.”* (Sphere Association et al., 2018). Therefore, in terms of equity, there is no reason why a person in a region with fewer cases is less deserving of aid than a person in a region with many cases. In other words: this formulation of the equity objectives rejects the idea that human suffering can simply be “added up” to determine where there is the highest need. This, obviously, influences the values seen in the model outcome. However, the simulation model is implemented in such a way that this definition could be changed easily if so desired by a problem-owner.

6.2. Validity of the Simulation Model for the Research Question

Apart from establishing that the simulation model behaves in a realistic manner, it also needs to be suitable for the purpose it will be used for: To discover the effect uncertain factors can have on the performance of treatment centre placement policies, and to determine what effect the incorporation of uncertainty-reducing decisions has on the performance of treatment centre placement policies for epidemics response. This section will address this issue by discussing the adequacy of the model scope and by arguing for the validity of the uncertain factors and their reduction.

First and foremost, the simulation model should contain all the concepts and behaviours that are relevant for the problem it studies. In this case, at the core, the problem is a spatio-temporal resource allocation problem under incomplete and uncertain information. Therefore, the model should include dynamic demand behaviour, a mechanism that allows for the placement of resources that fulfil demand, and a way to represent uncertain information and its reduction. In the simulation model, the dynamic demand is generated by the epidemiological model in the form of infected cases. The decision-module (representing the decision-maker) can place resources (ETCs) to meet the demand and prevent future cases from occurring. The implementation of uncertain information and the reduction of this uncertainty have already been discussed in detail in Chapter 4 and 5. Therefore, the key components relevant to the problem are represented in the simulation model.

However, it is necessary to question whether there are other concepts or behaviours which should be included. In many reports, newspaper articles and expert interviews, the distrust and political unwillingness were cited as one of the key elements hindering the response (see 9.4). This dimension is not present in the simulation model, the foremost reason being that these dynamics are very hard to translate into model components. What does that mean for the validity of the model itself? The aim of the study is to show what is *possible* in terms of policy performance given the defined resource allocation problem. Because the incorporation of uncertainty reduction through sequential decisions in resource allocation has not been studied before, it is more relevant to discover what types of model behaviour are possible without placing any exogenous constrictions on allocation decisions. Therefore the simulation model is said to be valid in terms of scope. Of course, the social and political dimension should be taken into consideration when translating the model results to actual policy advice. This will be discussed in detail in Chapter 9.4.

One of the core assumptions in this thesis was that the level of overall uncertainty experienced by the decision-maker could be used as the variable on which an adaptive policy is based. Which factors were uncertain and which ones were relevant for decision-making was determined in Chapter 3, based on literature on the response, reports by organisations active in the response, and news articles published during the crisis. Expert interviews were used to validate the relevance of these uncertain factors to allocation decisions. The uncertainty reduction as a result of allocation decisions is an essential part of the model behaviour, but very little academic literature is available on this issue. The assumptions on what regional information is or becomes available at an ETC are based on WHO guidelines for Ebola response World Health Organization and others (2014), which again were verified with expert interviews. However, none of the experts felt they could comment on the reduction of uncertainty resulting from an allocation problem, indicating that this was too complex a question. This resulted in the uncertainty reduction conceptualization being reliant on assumptions rather than empirical evidence. Given the model purpose this is acceptable (as the goal is to explore possible behaviours rather than mimic reality), but it is one of the weaker elements of the model.

6.3. Validity of the Research approach

Using MOEAs to optimize the policy function is a standard method in MORDM (Giuliani et al., 2015; Quinn et al., 2017). One limitation to the chosen radial basis function is that it cannot be parametrised to represent an fully-exploitative function. However, since an all-exploitative policy was run and incorporated in all experiments separately, this is seen as an acceptable limitation. Most importantly, the outcomes of the Borg MOEA showed that changing the type of decision made depending on the level of uncertainty has significant effect on the outcome objectives, and different objectives perform better under different policy functions. Further study and analysis showed that these different functions were sensible given the simulation model, the outputs they produced, and the objective(s) on which they were optimized. From this it can be concluded that the overall level of uncertainty as experienced by the decision-maker is a valid input factor to base an adaptive policy on.

7

Experiments

The simulation model as described in the previous section is used to run experiments to answer the remaining research questions:

5. Given the simulation model, what is the influence of system uncertainties on the performance of resource allocation policies?

6. Given the simulation model, which strategies for resource allocation decisions show robust performance?

This chapter describes the exact methods used to perform these experiments, and how the model is parametrised for each of the experiments. First, the EMA workbench, with which the majority of experiments is conducted, is described in brief. Next, the parametrisation of the simulation model for the experiments is outlined. The sub-question about the influence of system uncertainties is answered using exploratory modelling analysis, and the methods used for this are outlined. Finally, to identify candidate policies two many-objective evolutionary algorithms are discussed and the robustness measures on which the policies are evaluated are determined.

7.1. EMA workbench

In order to understand the influence of input uncertainties and different policies on the simulation model outcomes, the subject of sub-question 5, experiments are run using the EMA Workbench 2.0.3. The EMA Workbench is a Python library developed by Kwakkel (2017) which provides an interface to conduct exploratory modelling experiments, as well as analysis tools to be used on the results. The interface associates uncertain input factors, policy levers and outcomes to a simulation model and provides the functionality to run many simulations with different combinations of input factors and policies. The uncertain inputs of the simulation model defined in Chapter 5 are the number of initial cases in three districts: region 4 I_4, region 14 I_{14}, and region 15 I_{15}, the community transmission rate β_i, and the rate of travel r_{travel}. These factors are based on the uncertain factors identified in Chapter 3.

The type of policy inputs is dependent on which model is run - there is a model version in which the exploration versus exploitation ratio is constant throughout the simulation, in which case the policy is simply determined by providing the $exploration\ ratio$ as an input. In case the policy is based on the level of uncertainty as described in Section 5.3.3, the policy inputs required by the model are the variables of the radial basis functions: c_1, c_2, r_1, r_2, w.

The model outcomes are the scores of the five objectives (*Effectiveness, Cost per Death Prevented, Difference in Met Demand, Difference in Arrival Times, and Time until Containment*), which are calculated at the end of each simulation run as described in Appendix C.

7.1.1. Exploratory Analysis

In order to investigate the influence of uncertain factors on the outcome space of the model, exploratory modelling techniques were used. Experiments on the influence of the uncertain factors were carried out for two policies: a constant, fully exploitative policy, and a constant fully explorative policy. The all-exploitation policy is considered to be the best approximation of current practice, where all resources are allocated according to the humanitarian imperative of providing aid to those with the highest need, and no resources are spent on reducing uncertainty. The all-exploration policy is included to study any differences in the influence of uncertain factors that may occur on either side of the spectrum of possible policies, and only makes decisions aimed at reducing uncertainty.

The simulation model was run 2500 for scenarios for each policy (see appendix E for a validation of the ensemble size). In order to account for the stochastic elements within the model, replication were performed for each scenario, and the means of the outcomes of those replications were taken as the outcome values for that scenario. The number of necessary replications was determined by plotting the convergence of the mean for each of the objectives. At 50 replications, the strongest oscillations in the mean value were mostly settled, though in order to guarantee a stable outcome, up to 150 replications are needed for some of the outcomes. However, due to limited computational power, 50 replications were used in the experiments.

The scenarios were sampled over the uncertain input factors (using Latin Hypercube Sampling) with the following ranges:

$I_4 : 1 - 8$
$I_{14} : 20 - 35$
$I_{15} : 25 - 40$
$\beta_i : 0.1 - 0.5$
$r_{travel} : 0.04 - 0.1$

Where I_x represents the initial number of patients in region x. The range for β_i is based on Kucharski et al. (2015). Since no distinction is made between rural and urban regions in the simulation model, the range is chosen to represent all regions and extremes found by Kucharski et al. (2015) are not included. The range on the travel rates is based on the estimates made by Backer and Wallinga (2016).

In order to study how sensitive the all-exploitation policy is to the influence of the uncertain factors, and to see which factors have the most influence, several exploratory modelling techniques will be used:

- **Qualitative Visual Analysis:** By plotting objective scores against each other, patterns in model behaviour can be identified. This provides information of the spread of outcome values, the distribution of outcomes and how performance of objectives is related to each other.
- **Feature Scoring:** Feature scoring is a linear regression method implemented in the EMA Workbench which shows the individual influence of each uncertain factor (regressor) (Kwakkel, 2017). The level of correlation between the regressor and the outcome serve as a measure of sensitivity and therefore influence (Pianosi et al., 2016).
- **Dimensional Stacking:** Dimensional stacking is another method implemented in the EMA Workbench which determines the influence of uncertain factors, using feature scoring as described above. It visualizes the relation between the values of the most influential uncertain factors and model outcomes by creating a pivot table which shows the density of outcomes of interest for each combination of uncertain input factors.
- **PRIM**: The Patient Rule Induction Method, or PRIM, is an algorithm used to find regions in the input space that correspond to output regions of interest (i.e. with very good objective scores, or very bad ones that should be avoided). In this way, it provides insight into the model conditions which lead to certain results. PRIM works by searching for a box in the input space which contains both a high density of outcomes of interest (i.e. a large number of cases within the box are cases of interest) and has high coverage (i.e. a large number of the total cases of interest lies in the box) (Bryant & Lempert, 2010). PRIM creates rectangular boxes, which may lead to problems in finding an appropriate box if the cases of interest are not distributed in such as shape. In this case, PRIM results can be improved by applying Principal Component Analysis (PCA), which allows the rotation of the coordinate system such that cases of interest do form a rectangular

shape (Dalal, Han, Lempert, Jaycocks, & Hackbarth, 2013). The PCA transformation is available in the workbench.

7.2. Base Scenario

To serve as a point of reference throughout the experiments and further analysis, a base scenario was established. This scenario is initialized to reflect the state of the epidemic in Sierra Leone in June 2014.

For this base scenario, the uncertain factors are set to constant values. These are:

$I_4 : 3$
$I_{14} : 25$
$I_{15} : 31$
$\beta_i : 0.32$
$r_{travel} : 0.05$

The values for the initial number of cases are based on World Health Organization (2016), with region 4 representing Porto Loko, region 14 representing Kenema, and region 15 representing Kailahun. The community infection rate β_i is taken from Büyüktahtakın et al. (2018) and the travelling rate at 0.05 from Backer and Wallinga (2016).

This parametrization of the uncertain variables results in the simulation model closely following the real-case data (World Health Organization, 2016) for the the beginning of June 2014 to mid July 2014, which is assumed to be representative of the epidemic evolving in the absence of any response measures. The full parameterisation of the simulation model is outlined in Appendix D.

7.3. Direct Policy Search using MOEA

For the Direct Policy Search (DPS) the Borg Many-Objective Evolutionary Algorithm was initially selected to identify candidate policies. Borg uses a combination of different operators that mutate the input population. The amount of offspring each operator is allowed to produce is based on their performance during the run itself. Additionally, it uses a measure of progress called ϵ-progress which guarantees convergence and diversity of the solutions (Hadka & Reed, 2013). With this measure, the hyperspace of the objectives is divided into boxes with length ϵ (where ϵ can be different for each objective). Only solutions that improve performance and are not in an ϵ-box that is already occupied by another solution are counted as progress.

Using Borg, the parameters (c_1, c_2, r_1, r_2, w) of the policy function were optimized over the five model outcomes. Borg aims to find solutions that lie on the Pareto-optimal front. Since the policy function is normalized, the input ranges for the variables are as follows:

$c_1, c_2 : (-1, 1)$
$r_1, r_2 : [0, 1)$
$w : (0, 1)$

To set up Borg for with the simulation model, the following ϵ-values were chosen:

Effectiveness: 0.01
Time until Containment: 1
Difference in Met Demand: 0.01
Equity in arrival times: 10
Cost per Death Prevented: 50.

The ϵ-values for equity in arrival times and efficiency are relatively large, which in practice means scores on these objectives need to lie further apart in order for the difference to be considered significant. Since the values for these objectives are large (i.e. in the 1000s) this is justified, especially considering that modelling is aimed at finding patterns of behaviour rather than predicting exact values.

The Borg run was terminated at 36234 function evaluations after running for 92 hours on a laptop with a Intel Core i7-3630QM CPU, 2.4 MHz, and 6 GB RAM. At this point the run was terminated due to time constraints. The archive contained 718 solutions. The fact that Borg takes such a long time to converge is likely due to the complexity of the model (resulting in longer simulation time) and the stochasticity present in the model (in the "spontanteous news" function, in the random travelling mechanism, and in the selection of the decision type resulting from the policy function).

Since the Borg MOEA did not converge, the ϵ-NSGA2 MOEA which is implemented in the EMA Workbench was used, since this could be parallelized (as it is a steady-state MOEA) and run on a more powerful server. Additionally, this allowed for the replications to be incorporated into the MOEA. The ϵ-values were also doubled (0.02, 2, 0.02, 20, 100) to allow for quicker convergence. The ϵ-NSGA2 MOEA was first initialized for 25000 nfe with 150 replications, and was run on 3 servers to also allow for seed analysis. However, the high number of replications increased the computational burden to an unexpected and unworkable amount for the context of this thesis (i.e. requiring weeks to be completed). Due to limited time and access to computational resources the decision was therefore made to run the ϵ-NSGA2 MOEA with only 10 replications for 25000 nfe. Additionally, since the other servers where necessary to finish the policy experiments, no seed analysis was possible.

7.4. Robustness Testing

The policies generated by Borg and ϵ-NSGA2 were optimized on the base scenario described above. However, their performance should be evaluated on a wide range of scenarios in order to see how robust they are.

Robustness under deep uncertainty can be defined in a variety of ways which as a result can also influence the score and therefore ranking of policies in terms of robustness, as described in McPhail et al. (2018). They also provide a classification of various robustness measures, based on whether they measure absolute or relative performance, and how risk-adverse they are. Given that the problem owner in this context is the WHO, a global health security actor, who has to work with humanitarian organisations in a response, it is immediately clear that this is a risk-adverse decision-maker. The robustness measure(s) chosen to evaluate the policies generated by DPS should reflect this. Since different robustness measures can lead to different rankings of policies, it is also worthwhile to apply multiple definitions of robustness in order see if and why this also happens in this context.

Two robustness measures are chosen: The first is Undesirable Deviations, which is a relative regret based method. Given a scenario, regret-based robustness measures calculate the difference between the performance of the policy and the some reference performance, making it a relative measure of robustness. In Undesirable Deviations, the reference performance is the median value of the outcomes, and the regret values for the worst 50% of scenarios (50% percentile) are summed. The lower this sum, the more robust the policy. Undesirable Deviations is seen as a risk-adverse measure of robustness, though for instance a maxmin measure is even more risk-adverse McPhail et al. (2018). The choice is made for Undesirable Deviations since it takes into account a large amount of scenarios, which means that it can distinguish between policies which may have the same worst case scenario performance but different distributions for their overall performance.

The second robustness measure is Starr's domain criterion, which is an absolute robustness measure based on satisfaction of a threshold. This allows for the decision-maker to set certain thresholds or constraints on the objectives that should be met by a policy, and therefore also allows the decision-maker to determine the level of risk-adversion (McPhail et al., 2018). In the context of an epidemics response, it is for example realistic to assume that policies should have some acceptable level of effectiveness and that there is a limit to the amount of money organisations can justifiably spent on one crisis. Additionally, it allows for the policies found by DPS to be compared to the performance current practices. This is worthwhile as realistically, decision-makers can only be expected to change their strategies if new policies have been shown to at least meet current practice performances.

Each of the selected policies was run over 2500 scenarios with 50 replications, over the same uncertain inputs as described in Section 7.1.1. This provided the policy performance data on which the above robustness measures could be calculated.

8 Results

This chapter will describe the results obtained from the modelling experiments as described in Chapter 7. First, the model behaviour under two different policies (exploitative vs. explorative) is shown in terms of outcome distributions. The effect of the uncertain factors on the performance of these policies is analysed using exploratory modelling techniques. The runtime behaviour of the model and variations therein due to scenarios and policies is then outlined, providing additional insights into the model behaviour. The last section describes the performance of the policies that were found by the MOEA. This is done in three ways: their outcome distributions are compared, their runtime behaviour is outlined, and their performance on two robustness measures is calculated. The discussion and interpretation of these results is provided in Chapter 9.

8.1. Model Behaviour

In this section, the behaviour of the simulation model will be studied on three different aspects: First, the distribution of outcome scores for each of the objectives and any notable interaction between objectives is presented. Next, the influence of the uncertain factors on the outcomes will be studied using exploratory modelling techniques. Finally, the runtime behaviour of the model (i.e. what occurs during the simulation) is investigated. The analysis is carried out for two different policies: a fully exploitative policy, as well as a fully explorative policy. This is done to create a sense of the model behaviour at either end of the spectrum of possible policy inputs.

8.1.1. Outcome Distributions

The outcome distribution of each objective for the all-exploitation as well as the all-exploration policy are shown in Figure 8.1. In this section, the main features as well as interesting patterns visible in the outcomes will be discussed.

The all-exploitation policy's performance in terms of *Effectiveness* covers almost all possible values, with a peak of outcomes occurring at a score of around 80% to 85%, and another, lower peak at 30%. In contrast, the distribution of the outcomes for the all-exploration policy is much more narrow, with values ranging between 10% and 50%.

For *Time until Containment*, the majority of outcomes for the all-exploitation policy occur at 24 to 26 timesteps, with the distribution showing a tail towards lower scores which reaches outcomes just below 10 timesteps. Again, the distribution for the all-exploration policy is more narrow, ranging from values above 15 timesteps to the maximum of 26. Here the distribution shows two distinctive peaks, one occuring at 20-21 timesteps, and the other at 26.

For both policies, the outcome distributions for *Difference in Met Demand* have a similar shape, with one main peak, and a smaller peak to the right of it. For the all-exploration policy, this second peak is more pronounced, and the entire distribution is shifted to the right.

The outcome distributions for *Difference in Arrival Time* are most strikingly different: The all-exploitation policies shows a wide range of values (0 to about 1250) for this outcome, whereas the outcomes of the all-exploration policy all occur in a narrow range between 400 and 500. The fact that the all-exploration

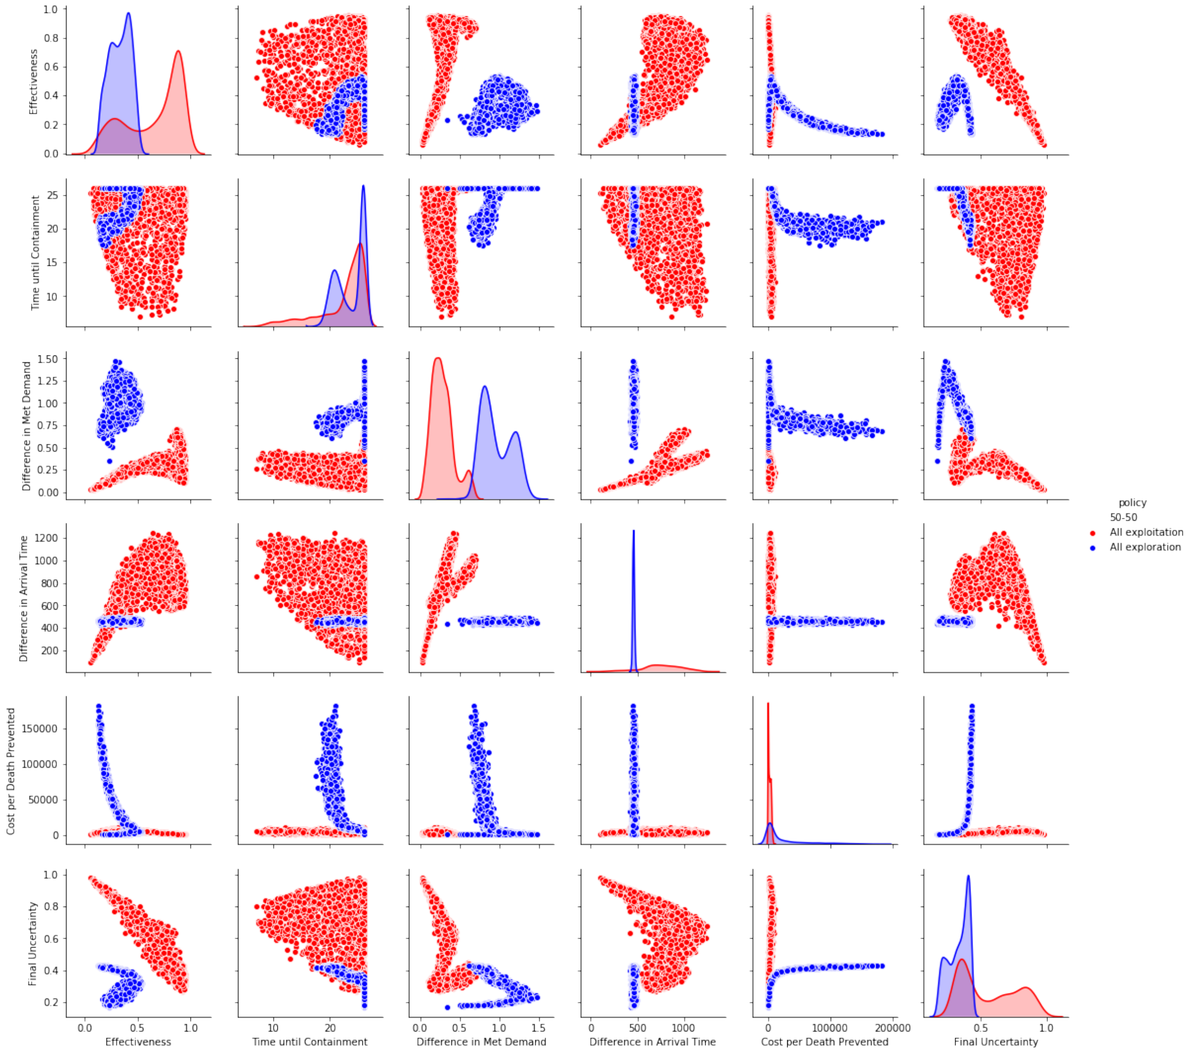

Figure 8.1: Scatterplot showing the objective outcome scores of the all-exploitation policy (red) and all-exploration policy (blue).

distribution has almost no variation in outcomes is easily explained when considering that the policy assigns resources based purely on the level of uncertainty present in each region. This will cause all the regions to be visited in a random but sequential order (i.e. a region is not visited for a second time until uncertainty in the other regions is reduced to a similar level) by this policy, limiting the variation in differences in arrival time that is possible.

For *Cost per Death Prevented*, the all-exploitation policy's distribution of outcomes is more narrow than that of the all-exploration policy, which has a considerable tail towards high cost outcomes.

For the additional outcome of *Final Uncertainty*, the distribution of the all-exploitation policy is wide and contains almost all possible outcomes, whereas the all-exploration policy's distribution is narrow and only contains low levels of final uncertainty. This is completely logical, since the latter policy is primarily focussed on reducing uncertainty. How the level of uncertainty evolves during runtime will be discussed in more detail in Section 8.1.3.

For the two policies, the interactions between various objectives are also noticeably different. The most striking interactions and differences will be discussed here. When plotting *Effectiveness* versus *Difference in Met Demand*, the all-exploitation policy relates the best outcomes in terms of Difference to Demand to either very low or very high Effectiveness scores. The combination of low Effectiveness and good performance in Difference in Met Demand can be understood as "equity in absence" (i.e. no one receives aid, and therefore everyone is equal), which is of course undesirable. For the all-exploration policy, no such interaction is visible.

Since the outcomes for Cost per Death Prevented from the all-exploration policy distort the outcomes of the all-exploration policy, the interaction of this objective was studied in a separate plot (available in Appendix G). This showed a strong negative correlation between the Cost per Death Prevented and Effectiveness scores, when Effectiveness scores were higher than 0.4. Below that, there is a positive, but less pronounced correlation. The all-exploitation policy also shows a clear negative linear correlation between *Effectiveness* and the level of *Final Uncertainty*. The all-exploration policy there is also an interaction between these two objectives, which takes an almost hyperbolic shape. The all-exploitation policy also shows some notable behaviour when the two equity objectives *Difference in Met Demand* and *Difference in Arrival Time*, are plotted against each other. Here the outcome distribution forks into two directions.

8.1.2. Influence of Uncertain Factors

In the previous section, it became clear that there are major differences in the outcome distributions of the two policies. This section will perform exploratory modelling analysis for each of the policies, in order to better understand how the uncertain factors (the transmission rate, the initial number of patient in three regions, and the travel rate) influence model behaviour [1].

For an initial side-by-side comparison, the outcomes of feature scoring analysis are shown side-by-side in Figure 8.2. For both policies, the transmission rate (labelled as beta_i) is the uncertain factor with the most impact. This is not unexpected, as the transmission rate is one of the dominant factors in determining disease dynamics (i.e. the number of future patients is influenced much more by the transmission rate than by the initial number of patients in a region). However, for the all-explorative policy, the transmission rate seems to be the only factor with a significant impact on the policy performance, except for the *Difference in Arrival Time* objective (which is expected, given the objective's insensitivity to input factors under an all-exploration policy, as explained earlier). For the exploitative policy, the influence is more spread for the *Time until Containment* and *Difference in Met Demand* objectives. As a first observation, we can establish that the fully exploitative policy is more sensitive to the uncertain factors than the fully explorative policy.

All-exploitative policy

Feature scoring showed that for *Effectiveness*, the transmission rate was the most influential uncertain factor. Visual inspection of the outcome distribution also showed two peaks - one at high (good) scores, and another at very poor scores. Using PRIM, first the desirable outcomes (scores above 80%) are studied. PRIM finds that these are strongly related to high transmission rates (ranging between 0.32

[1] The full analysis is available in a Jupyter Notebook called `ANALYSIS 2` at `https://github.com/edenbrok/thesis`

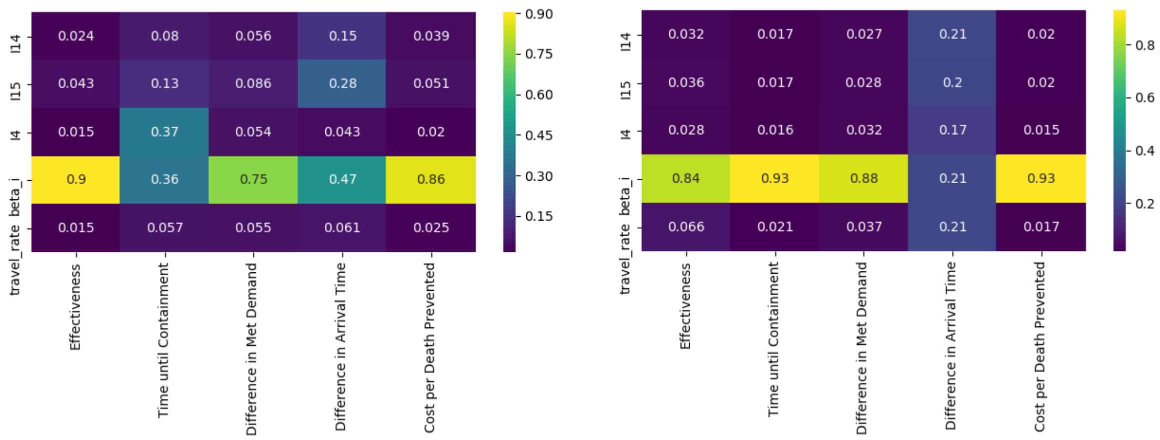

Figure 8.2: Feature scoring results for the all-exploitative policy (left) and the all-exploitative policy (right).

and 0.50). Similarly, poor performance (< 50%) is found to be related to low transmission rates (ranging beteen 0.10 and 0.22). A possible explanation for this influence is found in the spontaneous news mechanism. A high transmission rate will increase the number of cases quickly, increasing the chance of spontaneous news occurring very early on in the response. Since the all-exploitative policy is reliant on the spontaneous news mechanism in order to start its response (as otherwise it remains completely in the dark), receiving news early improves the quality of the response. This is confirmed by the strong correlation between Effectiveness and Final Level of Uncertainty which was identified in the previous section.

From the visual inspection of the outcomes, it was already clear that *Cost per Death Prevented* is strongly correlated with *Effectiveness* for the upper half of Effectiveness scores. For lower Effectiveness scores, however, a wider range of values is seen. Using PRIM on the outcomes with both low Effectiveness (< 0.4) and low Cost per Death Prevented (< 5000), three factors are found to be associated with these outcomes. A low transmission rate (0.10 to 0.23) is identified, as expected for the low Effectiveness score, as well as the initial number of patients in region 14 (20 to 30) and region 15 (25 to 34). For high costs (> 5000) and low Effectiveness, the association with a low transmission rate remains (0.10 to 0.19), but here the lowest numbers for initial cases are ruled out (with a range of 28 to 40 cases for region 15, and 22 to 35 cases for region 14).

In most scenarios, the all-exploitation policy reached containment (more than 70% of the current cases isolated) late in the response (i.e. only at timestep 25) or not at all. Therefore, we would like to know under what conditions fast containment (within 20 timesteps) occurs, as the outcome distributions do show it is possible. Using PRIM, the conditions for reaching good *Time until Containment* performance are found to correspond to a low number of initial cases (1 to 5) in region 4 combined with the lower half of possible transmission rate (0.10 to 0.32) values. Considering that for the objective 70% of the cases need to be isolated rather than 100%, it is likely that in these scenarios it was possible to reach this percentage by solely focussing on the hotspot regions 14 and 15.

Using dimensional stacking analysis, this explanation is confirmed, as the pivot plot in Figure 8.3 shows that these outcomes are also associated with a high number of initial cases in region 14 and 15.

In the previous section it was observed that for *Difference in Met Demand* there are two types of outcomes that have low (desirable) scores: Those associated with low effectiveness, which are essentially cases of ”equity in absence”, and good scores associated with high effectiveness. Since the factors associated with low effectiveness have already been established, we are mainly interested in identifying which conditions lead to the good *Difference in Met Demand* scores which are associated with an effective response. PRIM analysis determines that two factors are associated with these outcomes: a transmission rate range of 0.34 to 0.43, and values for the travel rate between 0.05 and 0.10, which is almost its entire range. The high transmission rate has earlier been established to be a condition for

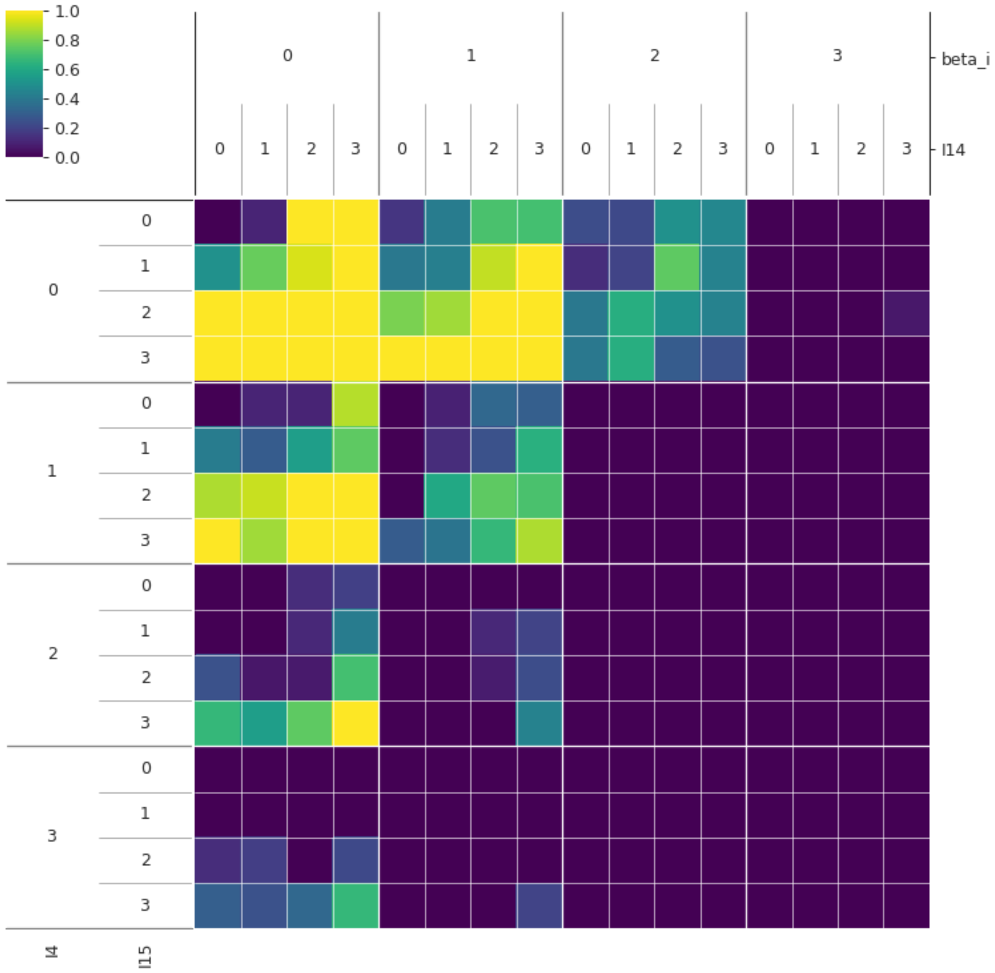

Figure 8.3: Dimensional Stacking pivot plot demonstrating that good outcomes in terms of Time until Containment are associated with low transmission rates, a low number of cases in region 4, and a high number of cases in region 14 and 15.

high *Effectiveness* scores. However, the very highest values are excluded here, which could be a result of the way in which the Difference in Met Demand objective is calculated. The objective score is based on the difference in the ratio of the total number of treated patients over the total number of infected cases. When the transmission rate is extremely high, for the hotspot regions (4, 14 and 15) this ratio will be dominated by the high number of infected cases that arises in the first few cases. As a result the difference between the low value of this ratio in the hotspot regions compared to the other regions is exacerbated and the resulting equity score is poor. This fits with the observation that the lowest travel rates are also excluded for good equity scores: when infected cases do not leave a hotspot region, this effect is also stronger.

For the *Difference in Arrival Time* objective the “best” scores are also clearly a case of equity in absence. What about the best performing half (scores below 750) that does occur with reasonable *Effectiveness* (> 0.75) scores? PRIM shows that these are associated with a high transmission rate (0.34 - 0.45) and the lowest values for the travel rate are excluded (with an associated range of 0.05 of 0.10). For poor scores in terms of *Difference in Arrival Time* (> 750) with high *Effectiveness*, the high transmission rate is included as well (0.29 - 0.50), and the lower half of possible values for the travel rate (0.04 - 0.08). Focussing solely on the worst outcomes in terms of *Difference in Arrival Time* (> 1000), by using dimensional stacking we can identify that these cases are associated with a relatively low transmission rate (0.2-0.3) and a high number of cases in region 15. This suggests that a particular range of transmission rate values (between 0.34 and 0.45) is a condition for good *Difference in Arrival Time* scores, but some other interaction must also be involved.

The effect of the travel rate seems counter-intuitive: one would expect high travel rates to have a negative effect on spatially defined equity since these cause the epidemic to spread faster. It is hypothesised that the negative influence of low travel rates on equity in arrival time is an artefact of the implementation of the simulation model: The *Difference in Arrival Time* objective uses the timestep at which a region first has more than 0 infections as the moment at which demand first occurs. However, an ETC is only placed when more than 1 case is observed (in order to prevent placement decisions based on 0.001 cases in a region where the epidemic is effectively under control). If region 15 is initialized with 32 cases, and the travel rate is set at 4%, region 11 will receive $\frac{32*0.04}{2} = 0.64$ cases, meaning it is not eligible for a placement decision even if the decision-maker has perfect information. With higher travel rates neighbouring regions receive enough cases to be eligible for placement decisions. Therefore, the way the objective is calculated causes the travel rate to have an unintended effect on the policy performance.

As for the *Effectiveness* objective, the policy is dependent on the spontaneous news mechanism in order to start placing the ETCs, which explains the lower bound on the transmission rate necessary to obtain good outcomes. Why the highest values for the transmission rate do not occur with desirable outcomes for *Difference in Arrival Time* is not directly clear, so this effect will be explored in the runtime analysis.

When comparing the relationships between objectives, it was noted that when *Difference in Met Demand* and *Difference in Arrival Time* are plotted against each other, their distribution shows a fork in the outcomes. In this section it has been established that the worst scores for *Difference in Met Demand* are associated with extremely high values for the transmission rate, whereas for *Difference in Arrival Time* the associated transmission rate range was between 0.2 and 0.3. This indicates that high scores for both *Difference in Met Demand* and *Difference in Arrival Time* are incompatible in terms of the scenarios in which they occur.

All-explorative policy

In terms of *Effectiveness*, the all-exploration policy was also most strongly influenced by the transmission rate. Again applying PRIM to the selection of best outcomes (scores higher than 40%), shows that the middle of available range for the transmission rate (0.25 to 0.39) is associated with these outcomes. For the poorest outcomes (scores lower than 40%), PRIM cannot identify a box, even when PCA is applied. Using dimensional stacking it becomes clear why: these outcomes are associated with either the lowest quadrant of values for the transmission rate, or the highest. This is unlike the clear division seen for the fully exploitative policy, and no explanation is immediately available. Since the all-exploration policy is not reliant on the spontaneous news mechanism for its decisions, this can be ruled out as a cause. The fact that low Effectiveness scores can be caused by either high or low values in the transmission rate does explain the the hyperbolic shape seen when plotting the *Final Uncertaint*y level against *Effectiveness*: for high transmission rates, uncertainty is reduced further due to the higher number of patients treated at ETCs. For lower transmission rates, the number of treated patients will automatically be lower.

The all-exploration policy has a very long tail in terms of outcomes with a high *Cost per Death Prevented*, which is also associated with low effectiveness. PRIM shows that the higher the cost, the smaller the associated range of low values for the transmission rate (with a range of 0.10 to 0.21 for costs above 50 000 dollars per life saved, which decreases to the range of 0.10 to 0.13 for costs above 100 000 dollar.) This is related to the poor scores for *Effectiveness* that occur at a low transmission rate. Why the highest transmission rates do not cause high costs (even though they are also associated with low effectiveness), will be explored in the runtime analysis.

For *Time until Containment*, the relationship between the input factors and the policy performance that becomes apparent through PRIM is straightforward. Late to no containment (> 24 timesteps) is associated with high transmission rates (0.29 to 0.50), better performance (< 24 timesteps) with low transmission rates (0.10 to 0.29). Since the all-exploration policy only starts placing ETCs later in the response (when no regions are hidden any more), high transmission rates will cause such a large number of patients that at this stage in the simulation run, containment is no longer possible with the random placement of small ETCs.

PRIM and Dimensional stacking relate the better scores in terms of *Difference in Met Demand* (< 1.0) for the all-exploration policy to the lower half of the possible transmission rate values (0.10 to 0.30). The highest transmission rates are not seen for the same reasons as set out in the analysis of the all-exploitation policy, but since the all-explorative policy is not dependent on the spontaneous news mechanism, good performance is also seen at the lower transmission rates.

8.1.3. Runtime Behaviour

How the level of uncertainty evolves during a simulation run is easily tracked, and the results of this are shown in Figure 8.12a for the all-exploration and all-exploitation policy (the third policy included in this plot will be discussed later in this chapter). From the plot, it is clear that the all-exploration policy manages to reduce uncertainty faster and more systematically than the all-exploitation policy, as would be expected. The wide range of possible uncertainty levels observed for the all-exploitation policy is caused by its dependence on the spontaneous news mechanism, as well as the uncertainty reduction it can achieve by the number of patients treated in an ETC (which also causes the variation in uncertainty seen for the all-exploration policy at higher timesteps). Notably, the all-exploitation policy also has a considerable number of scenarios where uncertainty is barely reduced. As seen earlier, these scenarios are associated with poor performance on almost all objectives (when taking into account the occurrence of “equity in absence”) and show the vulnerability of the all-exploitation policy.

In order to understand the runtime behaviour in depth, more factors need to be taken into account. Using data generated for each timestep during the simulation, the runtime behaviour of the model under different policies can be studied. For this, the actual number of cases per region are tracked, as well as the lower bound on the range of cases known by the decision maker. Additionally, the level of uncertainty in each region, and the type of decision made plus the chosen region(s) for that decision are kept track of. Visualizing all these aspects, the behaviour of the model during runtime can be studied using animations which show how the data changes over time. In this thesis only screenshots can be shown, but for each case a link and a QR code to the animation itself is also provided.

First, the `all-exploitation policy is run on the base scenario`. This run is shown in Figure 8.4. The all-exploitation policy receives spontaneous news only in timestep 4, though from both region 14 and 15 simultaneously. Here, the effect of the decision-maker only having access to an uncertain range around the ground truth becomes visible. Region 15 actually has the highest number of cases, but as the range known to the decision-maker is fitted randomly around the ground truth, the lower bound for Region 14 shows them a higher number of cases than for Region 15. As a result, an ex-

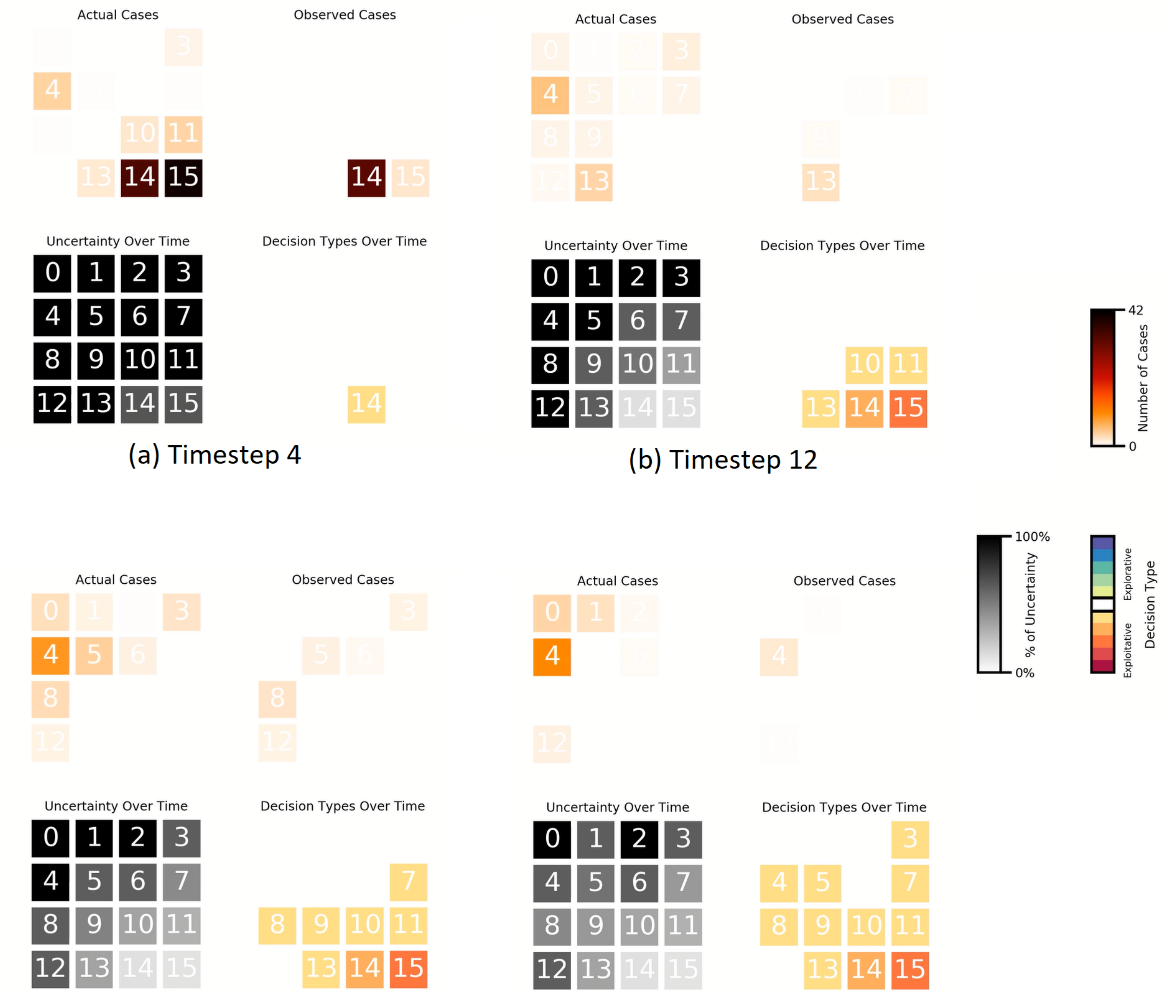

(a) Timestep 4 (b) Timestep 12 (c) Timestep 23 (d) Timestep 26

Figure 8.4: State of the simulation model under the all-exploitation policy and the base scenario at timesteps 4 (a), 12 (b), 23 (c) and 26 (d).

ploitative decision is made for Region 14. At timestep 12, a number of placement decisions have been made and ETCs are operational. Here, typical behaviour for the all-exploitation policy becomes visible: Through spontaneous news, uncertainty is reduced first in the hotspot-regions 14 and 15. As ETCs become operational, uncertainty is also reduced in neighbouring regions, and the response moves to those. This way, the situational awareness of the decision-maker unfolds over the neighbouring regions. The way uncertainty reduction spreads away from the hotspot regions can also be seen at timestep 23. However, this causes the response to remain centred around the initial hotspot regions, and in this particular run Region 4, which has had cases since the beginning, is only “discovered” in the very last timestep.

The runtime behaviour of the `all-exploration policy in the base scenario` is shown in Figure 8.5. At timestep 1, by pure chance it sends surveillance teams to regions 1, 14 and 15. As a result, the decision-maker immediately has awareness of the situation in two of the three hotspot regions. However, since the policy only makes explorative decisions, it cannot act on this information. By timestep 6 it has sent surveillance teams to all regions, and the policy will now start placing ETCs to further reduce uncertainty. Because it is placing the ETCs based on the level of uncertainty, the hotspot regions are “ignored” (see for example the plot at timestep 12), even though the decision-maker has a good overview of which regions have the most cases. By the end of the simulation run, it has placed

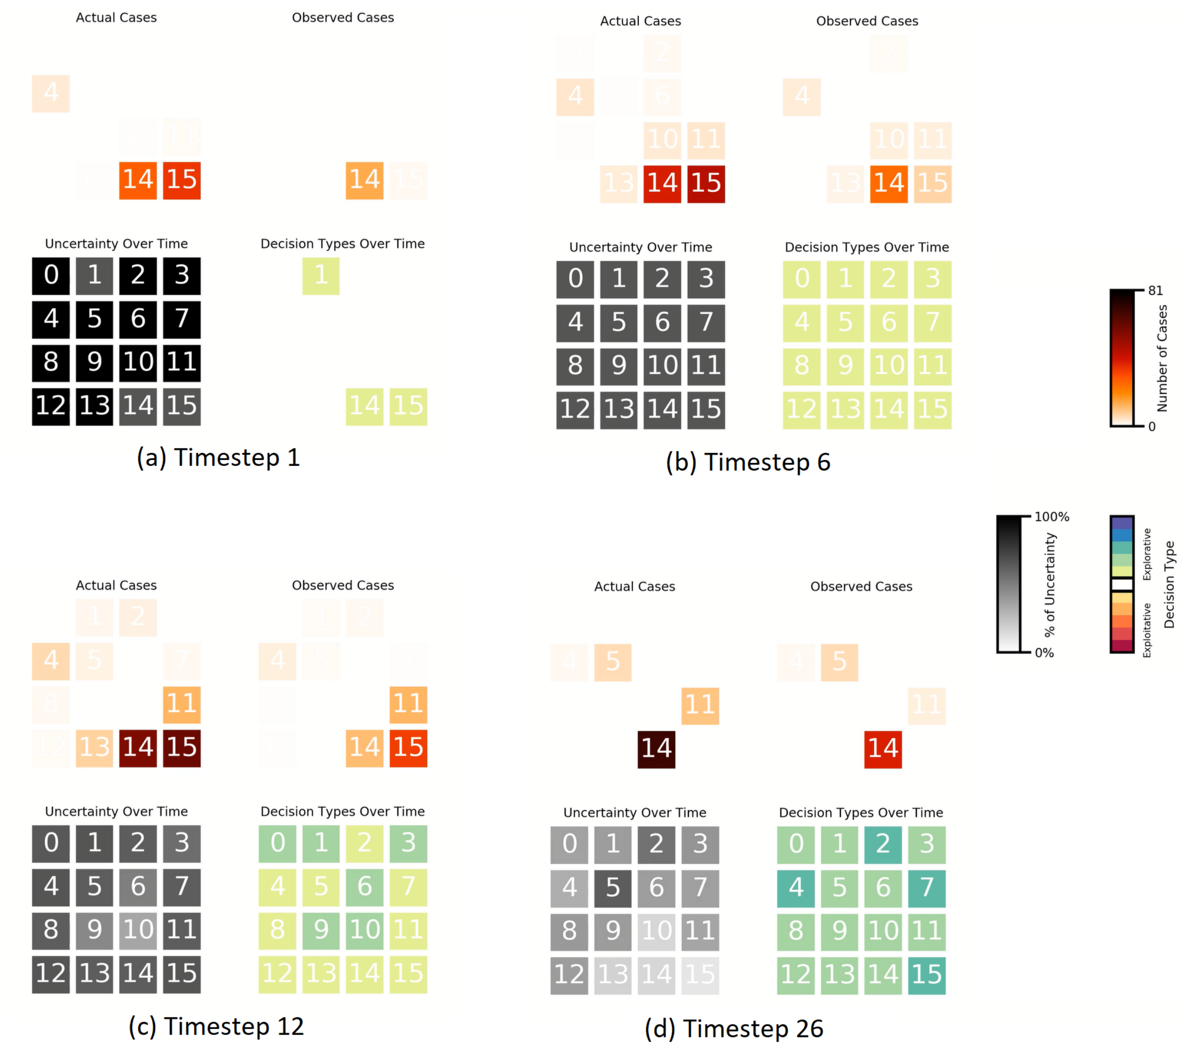

(a) Timestep 1 (b) Timestep 6 (c) Timestep 12 (d) Timestep 26

Figure 8.5: State of the simulation model under the all-exploration policy and the base scenario at timesteps 1 (a), 6 (b), 12 (c) and 26 (d).

ETCs in regions 4 and 15 by chance, but not in region 14.

Because the transmission rate was found to have the most influence on policy performance out of all the uncertain factors, animations are also generated with a low transmission rate of 0.1 and a high rate of 0.5.

Screenshots of the animation for the `all-exploitation policy with a low transmission rate` are shown in Figure 8.6. For this transmission rate, two types of model behaviour are seen: in one, the decision-maker never receives any spontaneous news and no placement decisions are made at all. In the other type of run, which is also the one shown here, the decision-maker does receive spontaneous news to act on. However, even in these scenarios performance in terms of *Effectiveness* is poor. Studying the runtime behaviour we can see why - due to the low transmission rate the number of cases starts to decline even before the response start. In the pictured response, an ETC is placed in region 15 during timestep 4, but at this point, majority of deaths have already occurred. This also explains the high *Costs per Death Prevented* associated with low *Effectiveness* scores: the ETCs that are built will treat only a small number of patients. As was identified in the exploratory analyis, a high number of initial cases softens this effect somewhat, which can lead to lower costs.

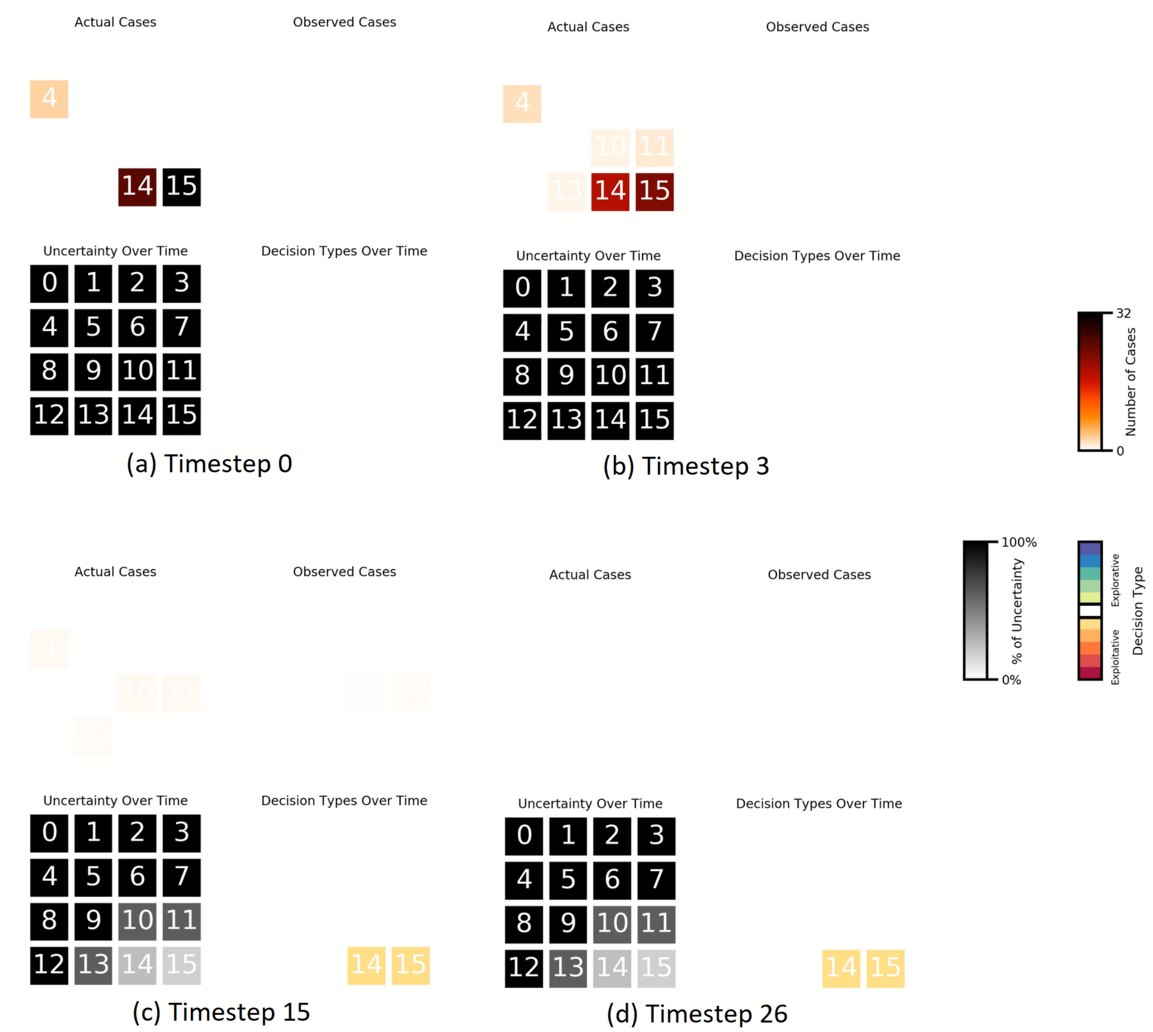

Figure 8.6: State of the simulation model under the all-exploitation policy and a low transmission rate (0.10) at timesteps 0 (a), 3 (b), 15 (c) and 26 (d).

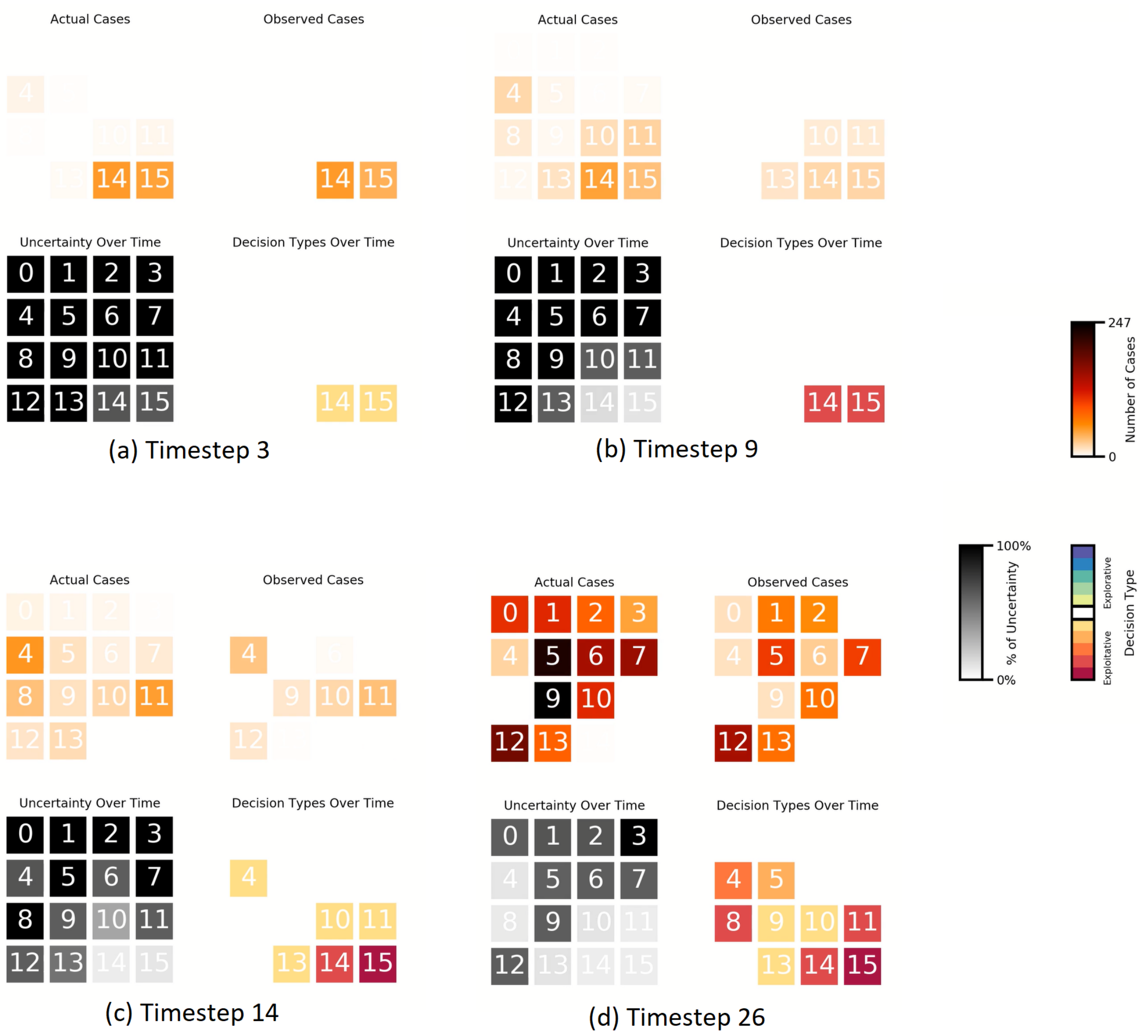

Figure 8.7: State of the simulation model under the all-exploitation policy and a high transmission rate (0.50) at timesteps 3 (a), 9 (b), 14 (c) and 26 (d).

The behaviour of the model under `the all-exploitation policy and a high transmission rate` is shown in Figure 8.7. Here, the decision-maker is aware of the situation in region 14 and 15 by timestep 3. Due to the placement of ETCs in those regions, the situation in the neighbouring region then also becomes known. However, due to the high number of cases in regions 14 and 15, the response remains focussed on the hotspot regions. This also explains the earlier observation that good scores for Difference in Arrival Time are not seen for the highest transmission rate values: these cause such an explosion of cases in the hotspot region that the placement of ETCs remains focussed there, even though the decision-maker is aware of the situation in neighbouring regions. As a result, the regions which are not hotspots have to wait longer for aid. With the high transmission rate, the decision-maker also receives spontaneous news from region 4 in this particular run. Though a high number of cases still remains at the end of the simulation run, the policy performs well in terms of *Effectiveness*. This is because without a response, the high transmission rate would have lead to an exponential growth of cases, a large proportion of which is prevented by the response.

For the `all-exploration policy under the low transmission rate`, we see the same effect of the epidemic decreasing in severity even before the first ETC is placed, which leads to poor *Effectiveness* scores. However, the all-exploration policy avoids the worst case scenario of the all-exploitation policy that never makes any placement decisions, because after surveillance teams have been sent to all regions (at timestep 6), it starts placing ETCs in all regions. Though these are

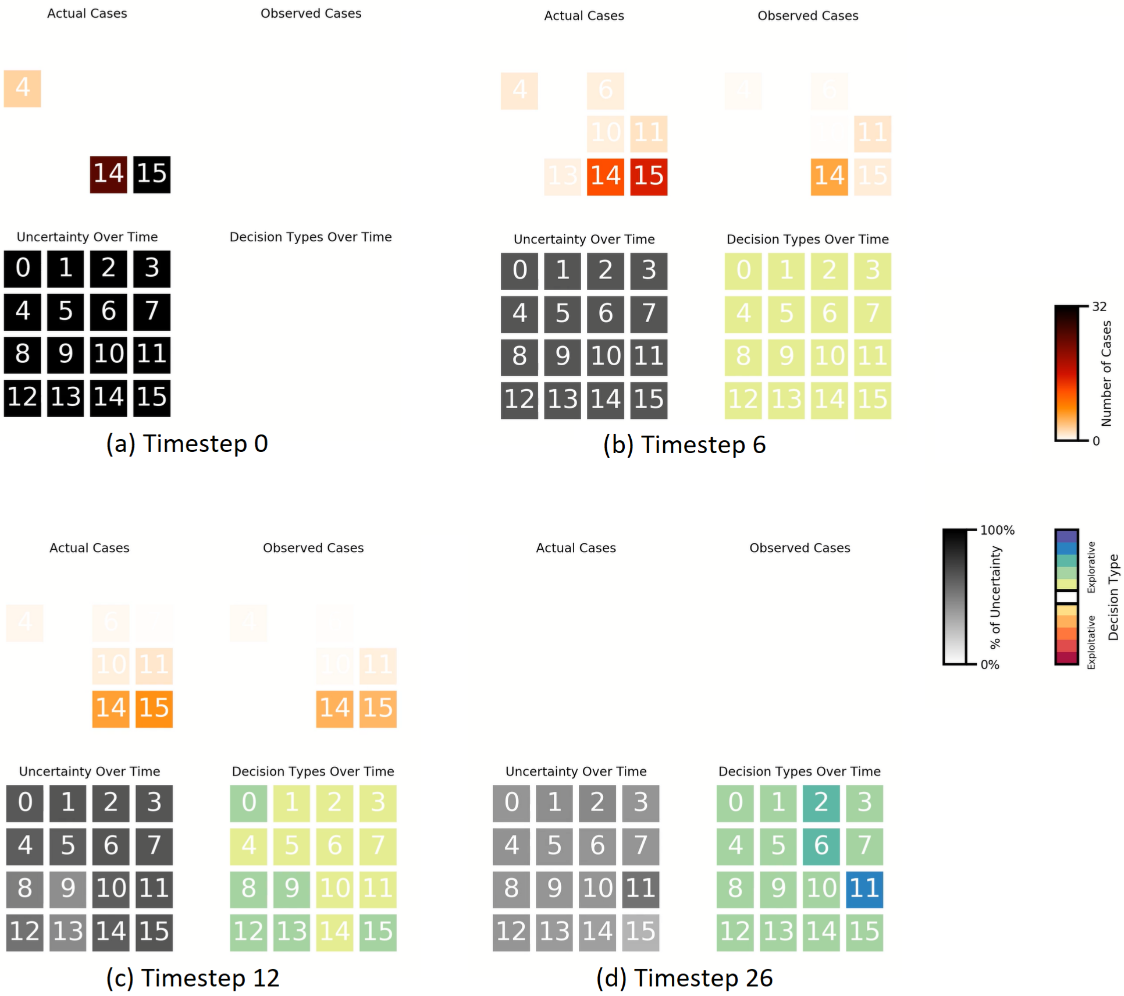

Figure 8.8: State of the simulation model under the all-exploration policy and a low transmission rate (0.10) at timesteps 0 (a), 6 (b), 12 (c) and 26 (d).

not placed based on the number of cases, inevitably some of them are placed in the hotspot regions.

With an understanding of how this policy operates, the runtime behaviour of the model with a `high transmission rate under the all-exploration` policy bears little explanation. The situation at the last timestep is shown in Figure 8.9. Note the scale of the axis indicating the number of cases - over 7000 cases are reached in some regions. In the exploratory modelling analysis of the all-exploration policy it was noted that high costs occur at low *Effectiveness*, but only for the lower values of the transmission rate (even though low effectiveness was found to be caused by both very high and very low transmission rates). By studying the runtime behaviour we can understand why: Under the low transmission rate, the epidemic has declined already when the first ETC has become operational, meaning that effectiveness is low because most deaths have already occurred. Because there are only a small number of cases left, costs per case treated are proportionally very high. With the high transmission rate, by the time the first ETCs become operational the epidemic has already grown completely out of control, but given the high number of patients costs are proportionally lower.

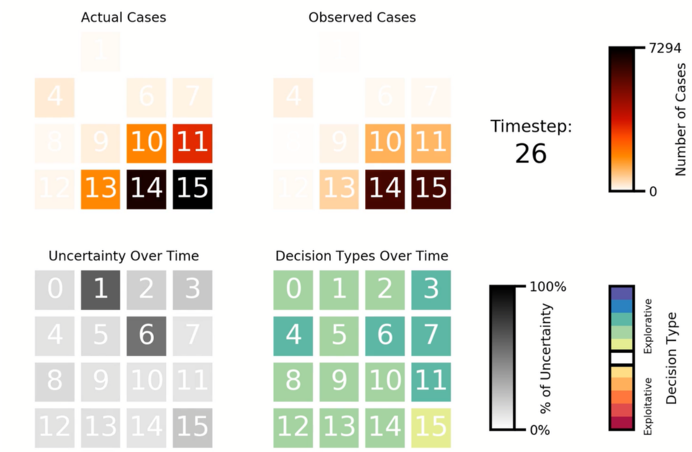

Figure 8.9: State of the simulation model at timestep 26 under the all-exploration policy with a transmission rate of 0.5.

8.2. Policy Performance

This section discusses the performance of the policies that were optimized for the base scenario using MOEAs. A full discussion of the MOEA results can be found in Appendix F. From the solution set, five policies were selected, each representing optimal performance over one of the objectives. They are shown in Figure 8.10. The rest of this section is structured as follows: First, their output distribution over a wide variety of scenarios is considered. Next, the runtime behaviour of these policies is analysed in the same way as was done for the all-exploitation and all-exploration policy above. The last section determines their scores for the two robustness measures [2].

For clarity, the five selected policies are classified according to the behaviour of their policy function at high uncertainty (i.e. at the beginning of a simulation run). Policies 390 and 661 are exploitative policies, policy 148 is a mixed policy, and policies 141 and 185 are explorative. From here on, they will be referred to with their classification followed by their number (i.e. exploitative-390). The fully exploitative policy, the fully explorative policy and a “random” policy (which chooses between explorative and exploitative actions with equal probability and is classified as a mixed policy) are also included in some figures to act as a reference.

8.2.1. Policy Behaviour per Classification

Outcome distributions of the different policies are shown per groups and per objective in Figure 8.11 (full scatterplots of the outcomes for each policy can be found in Appendix G). For *Effectiveness*, the exploitative policies have very similar distributions. The same general shape is also seen in the mixed policies, although here the worst outcomes are avoided and the peaks associated with good performance are higher. These differences become even more apparent with the explorative policies. However, for these policies, the peak with the best performance occurs at a lower *Effectiveness* score than for the exploitative and mixed policies. The only policy with a truly different outcome distribution shape is the fully explorative policy, suggesting that shifting to exploitative decisions at lower uncertainty has a significant impact on policy performance.

For *Cost per Death Prevented*, all policies follow an exponential distribution, though the exploitative policies show a second, much smaller peak. Given the strong correlation between effectiveness and cost observed earlier, this second peak most likely corresponds to the slight peak the exploitative policies show at low *Effectiveness* scores. It can also be observed that the more explorative a policy is, the longer the tail towards higher costs is. Though not included in these plots, exploitation-390, exploitation-661 and mixed-148 all had outcomes with outlier values ranging from the order of 10^7, whereas the all-exploitation policy did not have these outliers. How these extreme outliers are caused could not be identified.

The distributions for *Time until Containment* for the exploitative policies all peak around 25 timesteps, and have a tail towards lower outcomes, down to 5 timesteps. For the exploitation-390 policy, the peak is more narrow and to the right of that of the all-exploitation policy, whereas for exploitation-661 the exact opposite is the case. The mixed policies both have wider peaks than the exploitative policies, with more outcomes at lower values. Notably, the random policy shows two slight peaks. The all-exploration policy, as well as the explorative-185 policy also have two distinctive peaks, which are much more pronounced. However, the explorative-141 policy does not show this behaviour.

Interestingly, for *Difference in Met Demand*, all policies follow the same general shape of a major peak and a minor peak to the right of it (at poorer scores). It also seems like the more explorative a policy is, the more the entire distribution is shifted towards the right. The major peaks for the exploitative policies, as well as the mixed-148 policy are also wider and their shape suggests they consist of two peaks laying close together, whereas for the other policies this is not the case.

For the *Difference in Arrival Time* objective, the policy classifications clearly do not correspond to similarity in outcomes. The cause of the narrow peak of the all-exploration policy has been discussed in Section 8.1.1, but the exploitation-390 policy and the explorative-185 policy also have narrow distri-

[2] The plots used in this section were generated in the `ROBUSTNESS ANALYSIS 2` jupyter notebook which is available at `https://github.com/edenbrok/thesis`

Policy Name	Policy Function	Objective Score: Effectiveness	Cost per Death Prevented	Difference in Met Demand	Difference in Arrival Time	Time Until Containment
all-exploitation	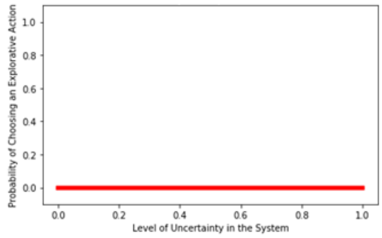	0.815579	2136.118	0.379401	991.75	26
exploitative-390	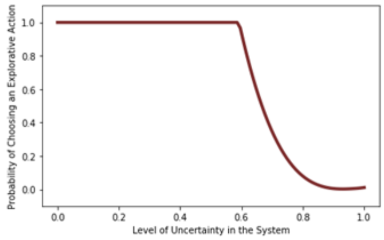	0.883009	3121.1938	0.32228	803.9375	**5**
exploitative-661	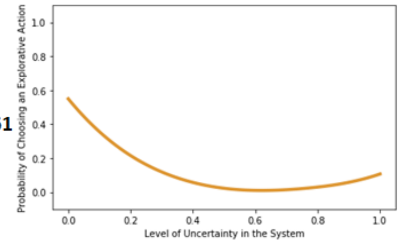	0.851158	**1055.8085**	0.538239	1270	6
mixed-148	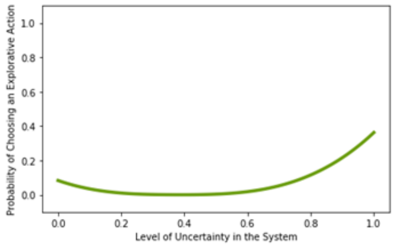	**0.881327**	2123.2283	0.457683	1117.75	6
explorative-141	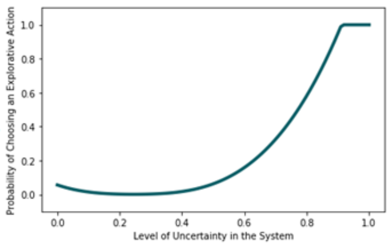	0.751473	3274.8133	**0.033334**	399.4375	25
explorative-185	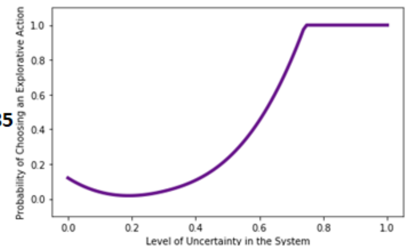	0.724772	3649.6277	0.496914	**333**	20

Figure 8.10: Selected policies from the Borg solutions with their associated objective scores. The objective on which the policy had an optimal score is shown in bold. The all-exploitation policy is shown as a reference.

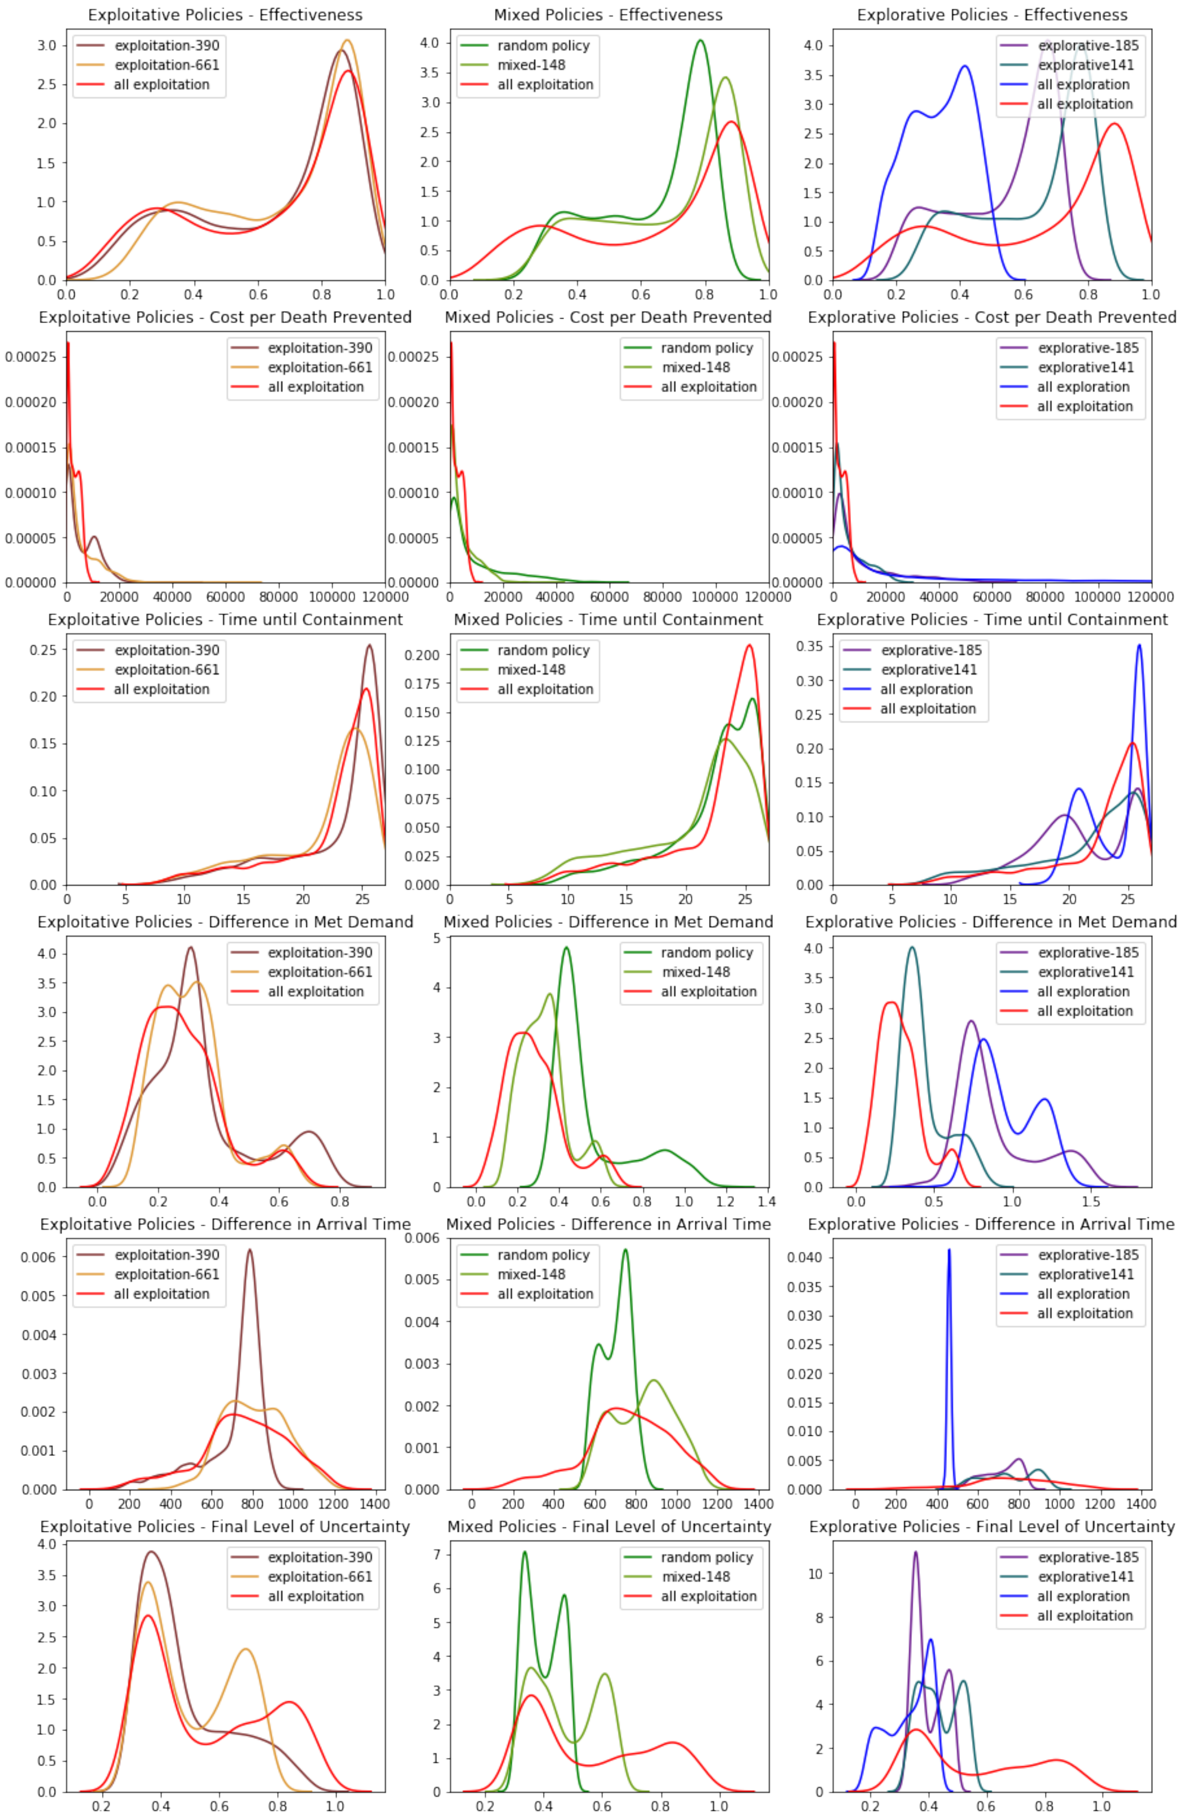

Figure 8.11: Outcome distributions for all objectives, including the additional outcome of the level of uncertainty at the end of a run. The all-exploitative policy (in red) is included in all plots for reference.

butions with just one peak. All other policies have wider distributions, in some cases with two peaks (exploitation-661, both of the mixed policies, and exploration-141).

For the *Final Level of Uncertainty*, two general observations can be made: Firstly, the more explorative a policy, the more narrow its distribution. Secondly, all distributions show two peaks (in the case of the exploitation-390 policy a distinct second peak is not visible, though it does have a pronounced shoulder). For the explorative policies, the leftmost peak is the most pronounced (at lower uncertainty levels), whereas for the mixed policies the two peaks are much closer in height. For the explorative policies, no such pattern is visible, in fact, for the all-exploration policy the rightmost peak is the highest, for the explorative-141 policy they are both the same height, and for the explorative-185 the leftmost is the highest.

How the level of uncertainty develops during runtime, leading to these final uncertainty level distributions, will be investigated in the first part of the next section.

8.2.2. Runtime Behaviour

Uncertainty Reduction over Time

The total level of uncertainty experienced by the decision maker is shown in Figure 8.12b for the exploitative policies, in Figure 8.12c for the mixed policy and in Figure 8.12d for the explorative policies. Like the all-exploitation policy, the exploitation-390 and exploitation-661 policies demonstrate the widest range of total uncertainty over all the scenarios. The range of the exploitation-390 policy contains a significantly larger amount of scenarios in which uncertainty remains high throughout the simulation run compared to exploitation-661. This can be explained by comparing their policy functions: Though the exploitation-390 policy moves to explorative actions at higher levels of uncertainty than the exploitation-661 policy, at full uncertainty the latter takes explorative actions with a chance of 10%, whereas the exploitation-390 policy only does so with a chance of 1%. This makes it reliant on the spontaneous news mechanism, and like the all-exploitation policy, increases its risk of being immobilized by a lack of information.

Like the random policy, the mixed-148 policy's range fans out later in simulation time. It ends up with a wider range of possible final levels of uncertainty, which is expected as it is less “mixed” than the random policy, starting with a 36% chance of taking explorative actions (as opposed to 50% for the random policy), and this chance is reduced to 0% when the perceived uncertainty is at 50%.

The two explorative policies show similar progressions of uncertainty over time, with the main difference resulting from the level of uncertainty at which each policy moves to mixed/exploitative decisions. Like the all-exploration policy, the exploration-185 policy remains fully explorative at high uncertainty levels. This results in the small “plateau” which occurs for both these policies at an uncertainty level 66.7%. This is the level of uncertainty at which all regions are discovered (i.e. no longer hidden) but at full uncertainty. As the explorative-141 policy is already taking exploitative decisions at this level of uncertainty, it moves past this plateau as it is no longer basing decisions on the highest number of level of uncertainty.

No policy reaches uncertainty levels below 30%, except the fully-exploitative policy, which lowest reached level is 17%. Since uncertainty on the number of patients in a region cannot be reduced below 20% due to case hiding (see Chapter 4 and Appenidx B), the lowest possible value for total system uncertainty is 6.7%.

Based on these results, it can be observed that the model behaviour in terms of uncertainty and its reduction over time seems to be largely determined by the shape of the policy function at high values uncertainty levels: Exploitative policy 390 becomes fully explorative at lower levels of uncertainty, yet its behaviour in terms of uncertainty is very similar to that of the fully explorative policy. Likewise, the mixed 148 policy is fully exploitative at lower uncertainty, but its uncertainty reduction resembles that of the random policy, and though its final distribution is wider, the final outcomes are distributed over two peaks of similar height, as is the case for the random policy.

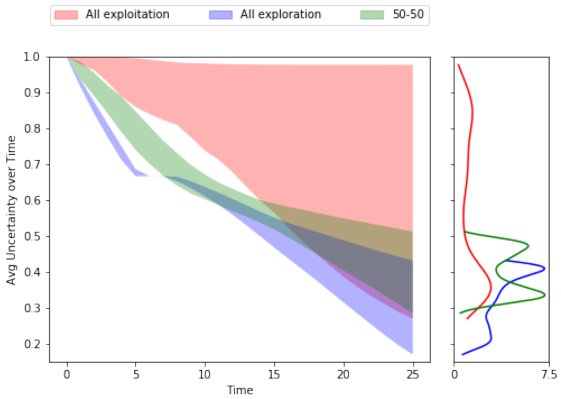

(a) Network 1

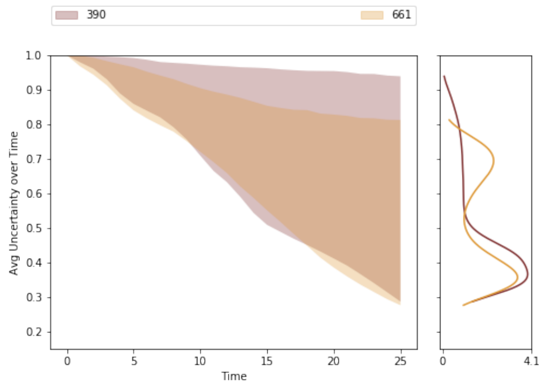

(b) Network 2

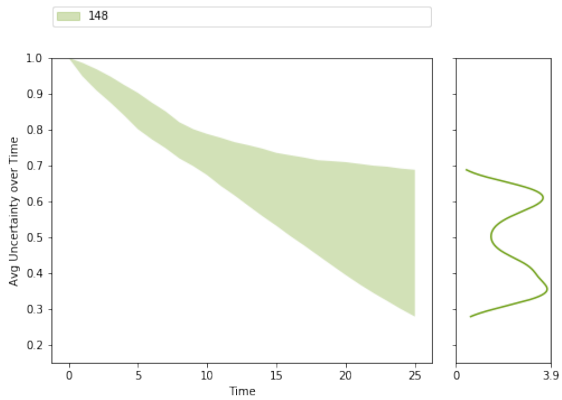

(c) Network 3

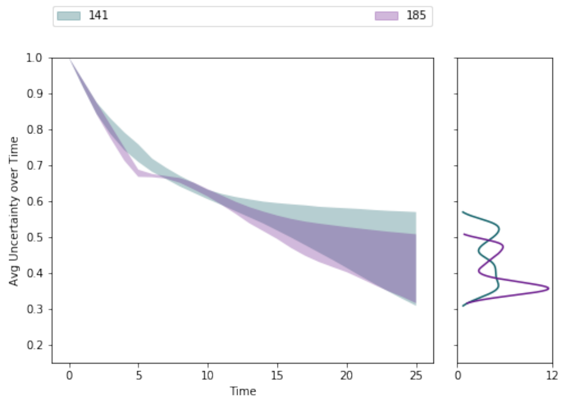

(d) Network 4

Figure 8.12: Plots showing the development of the total level of uncertainty as perceived by the decision maker over time. The plot on the left shows the envelope of all outcomes in the 2500 scenarios, the plot on the right the distribution of final uncertainty levels over all these scenarios.

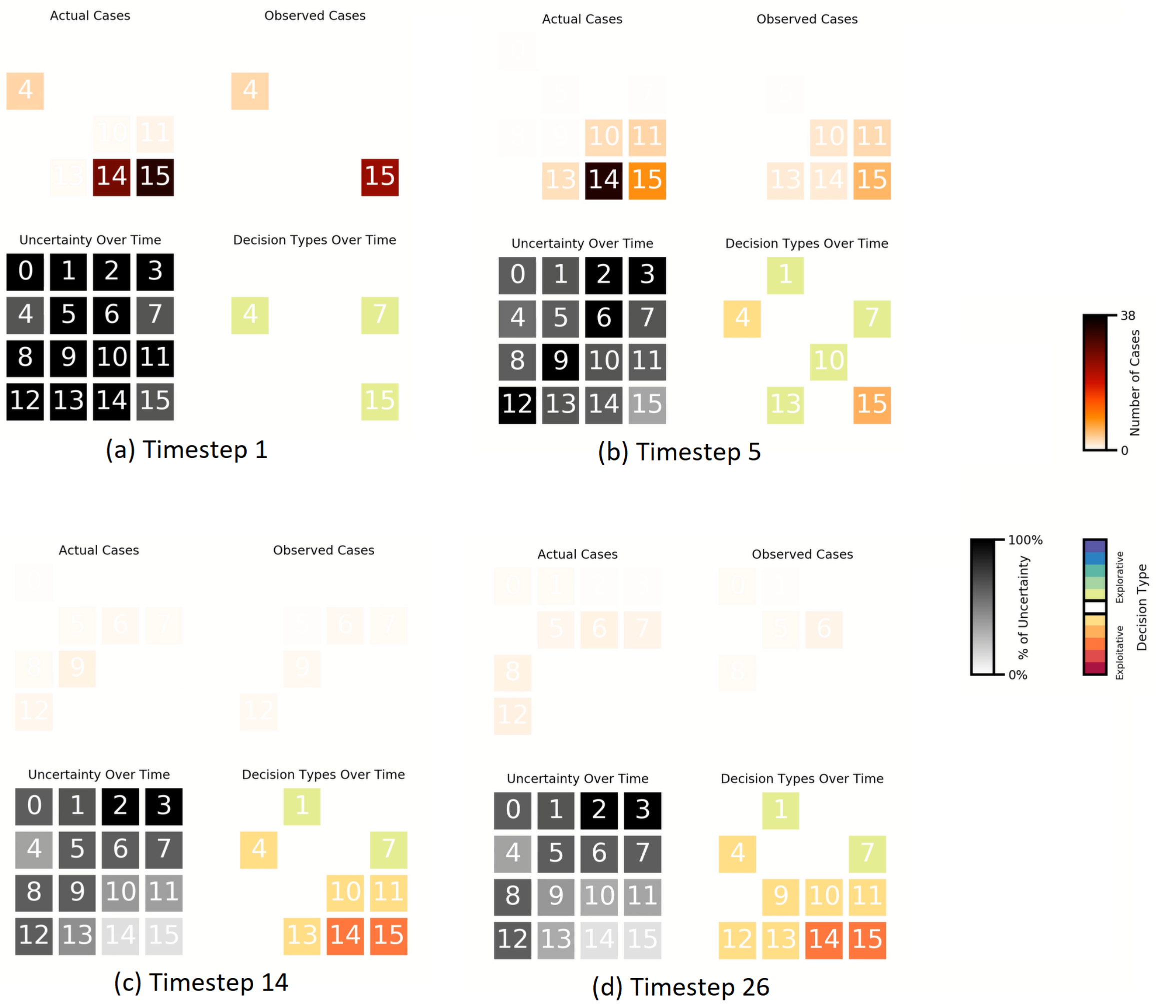

Figure 8.13: State of the simulation model under the exploitation-661 policy at timesteps 1 (a), 5 (b), 14 (c) and 26 (d).

Runtime Behaviour

The runtime behaviour of the five policies is studied in the context of the base scenario (as defined in Chapter 7.2). For brevity, only one policy from each classification is described here. The descriptions of the runtime behaviour of the other policies can be found in Appendix H.

The `runtime behaviour of the exploitation-661 policy` is shown in Figure 8.13. Its first decision is explorative, and it makes the very lucky decision to send surveillance teams to two of the hotspot regions. This means that by timestep 5, it has already placed three ETCs. However, due to the uncertainty around its observations on the number of cases, the severity of the epidemic in region 14 has gone unnoticed. By timestep 14 this has been corrected, and ETCs have been placed in all the hotspot regions as well as in some of the neighbouring regions. The overall level of uncertainty has now decreased to a point where the policy is fully exploitative, causing the response to remain almost entirely static until the end of the simulation run. It should be noted that, had region 4 not been discovered by pure luck early on in the response, it would have likely gone completely unnoticed. As it is, even at the end of the response regions 2 and 3 remain completely hidden from the decision-maker.

In Figure 8.14 the `runtime behaviour of the mixed-148 policy` is shown. Within the first three timesteps, it has taken one explorative decision, and two exploitative decisions, of which only the second resulted in a placement decision since the decision-maker had received spontaneous news from region 14. At timestep 6, it is still alternating between decision types, and the decision-maker has gained awareness of the situation in region 15 thanks to its presence in region 14. By timestep 14,

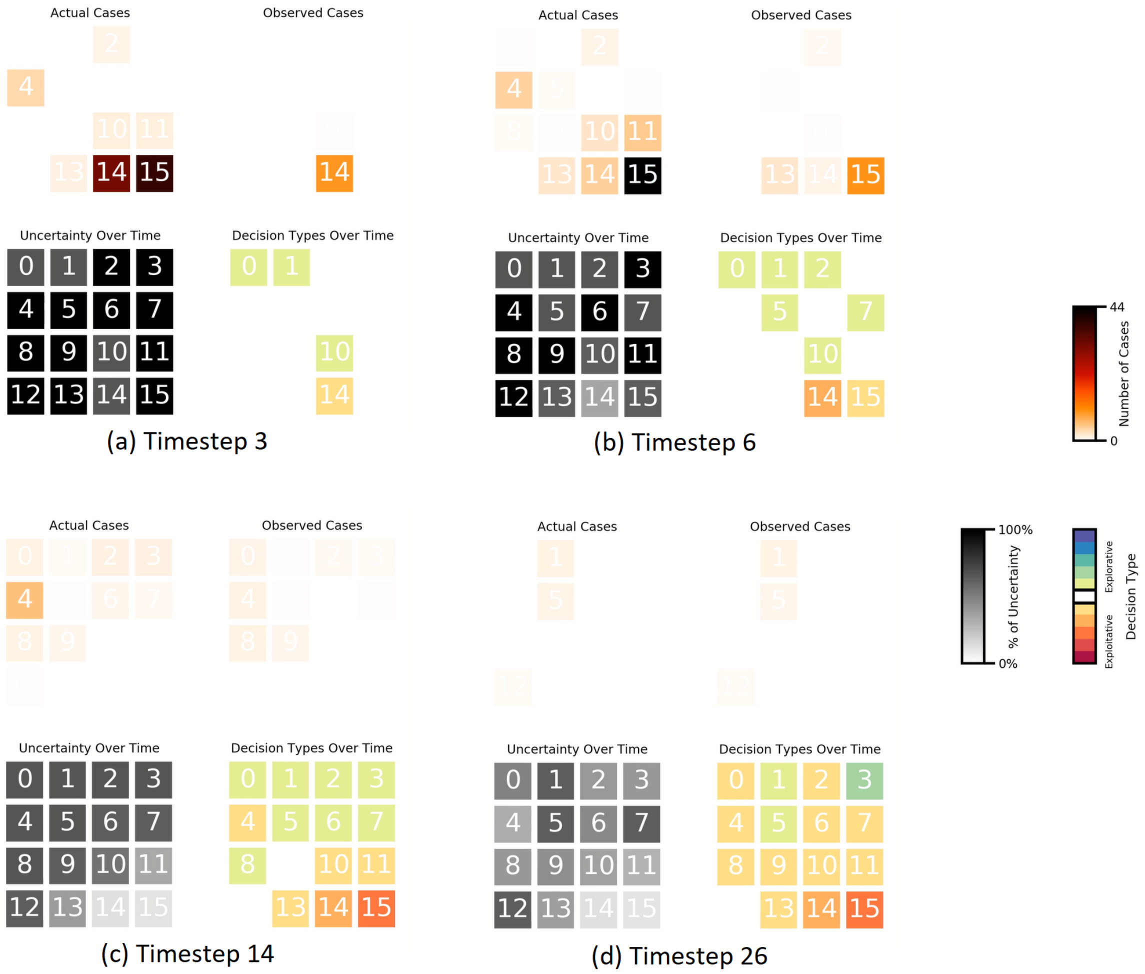

Figure 8.14: State of the simulation model under the mixed-148 policy at timesteps 3 (a), 6 (b), 14 (c) and 26 (d).

none of the regions remain hidden, either by explorative decisions or due to the presence of an ETC in a neighbouring region. Until the end of the simulation run, the policy then takes primarily exploitative actions. At the end of the run, 12 different regions received an ETC as a result of an exploitative decision, which is even more than seen for the all-exploitation policy in Section 8.1.

The `runtime behaviour of the exploration-141 policy` is shown in Figure 8.15. At timestep 3, the two explorative actions have reduced the overall level of uncertainty enough for the policy to have moved to a mixed stadium. At this step, it takes an exploitative action, and since region 15 is no longer hidden, it can make a placement decision for that region. In timestep 9 we can see that the policy has remained mixed. None of the regions are hidden any more, and three ETCs have been placed with exploitative decisions. At timestep 18, we can see that most decisions are now exploitative ones, but explorative decisions still occur. Between timestep 18 and 26, only exploitative decisions are made, and in several cases the decision maker chooses to do nothing due to a lack of (observed) cases. Though at the end of the simulation the epidemic is contained and the decision-maker has a good situational overview, it should be noted that due to the emphasis on explorative decisions at the beginning of the response, only a single ETC was operational before timestep 9. This illustrates the trade-off between explorative and exploitative actions and explains why the explorative policies never reach the high Effectiveness scores of the other policy types.

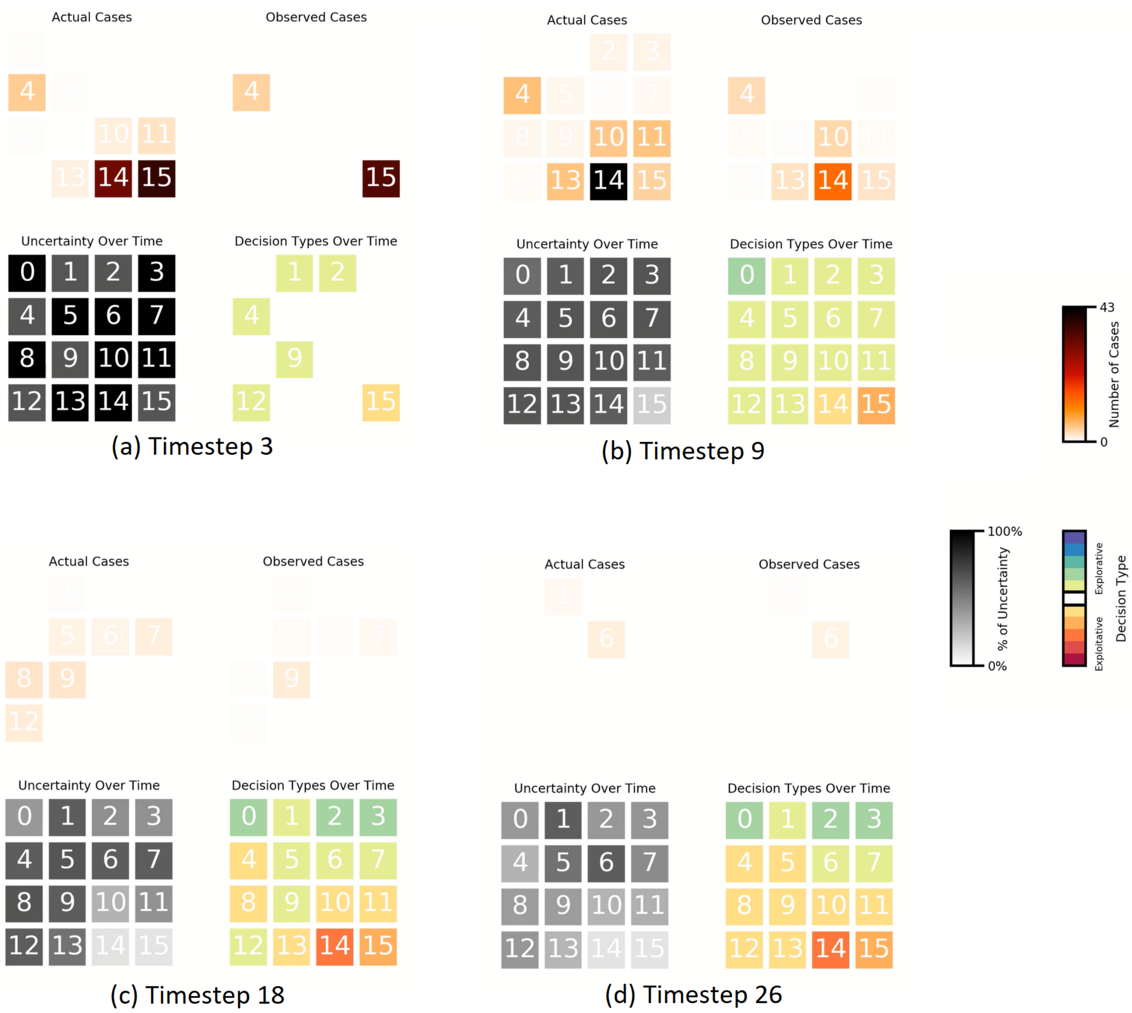

Figure 8.15: State of the simulation model under the exploration-141 policy at timesteps 3 (a), 9 (b), 18 (c) and 26 (d).

8.2.3. Robustness Scores

Undesirable Deviations

For all five policies, the undesirable deviations scores on each of the objectives were calculated and then normalized. The results are shown in the parallel coordinates plot in Figure 8.16, which also includes the all-exploitation policy which serves as a reference for current practice. The best scoring policy on a particular objective is assigned the value of 0, the worst a score of 1.

For the *Effectiveness* objective, the difference in robustness is clearly determined by the type of policy: the exploitative and mixed policies, with wider outcome distributions have high scores, and the two explorative policies score best. The opposite is visible for the *Cost per Death Prevented* objective. Even though the exploitative and mixed policies had a few outliers, the explorative policies, and in particular the explorative-185 policy score much worse because they have a much longer tail towards high costs. This policy also scores the worst in *Difference in Met Demand*, where it had a wide distribution and worse scores than the other policies in general. Here the mixed-148 policy scores best, followed by the exploitation-661 and all exploitation policy. The exploration-141 policy and exploitation-390 policy similar scores, with the latter performing slightly worse.

For *Difference in Arrival Time* the policies which have narrow distributions (see 8.2.1), unsurprisingly, have the best score. Again the policy classification does not serve to predict how robust the policy will be: the best and worst performing policies are exploitative, and the explorative and mixed policies make up the middle scores. Classification also provides no guarantees for performance in robustness on *Time until Containment*. The all-exploitation policy performs best, followed closely by the two explorative policies, then the mixed policy, and the other exploitative policies perform the worst, particularly the exploitation-661 policy. This is a result of those policies having long tails towards good outcomes but a peak at the worst outcomes.

Starr's Domain Criterion

The scores of the five policies on the Starr's Domain Criterion robustness measure are shown in Figure 8.17. Since in measure the median of the all-exploitative policy is used as a reference, the all-exploitative policy scores exactly 50% on all objectives (as per its definition, 50% of its outcomes will lie above the mean). Therefore, any scores above this red line indicate a better performance than that of the all-exploitative policy.

For *Effectiveness*, the scores of the exploitative and mixed policies lie very close together (within a range of 0.01 around 0.5). Unsurprisingly, the explorative policies perform much worse, which is in line with the earlier observation that their outcome distributions had their peaks at lower *Effectiveness* scores than the exploitative policies. For the *Cost per Death Prevented* objective, however, all policies except the exploitation-390 policy score better than the reference reference policy. This means all these policies perform better than the all-exploitation policy in more than half of the scenarios. A different picture is visible for *Difference in Met Demand*, where none of the policies can meet the performance of the all-exploitation policy, and the explorative policies perform particularly poorly. This is expected given the observation that their distributions for the objective were shifted to the right considerably when compared to the other policies. In terms of *Difference in Arrival Time*, the two explorative policies do perform better than the all-exploitation policy. This also makes it notable that the mixed-148 policy perform the worst in terms of robustness on this objective. Studying its distribution of outcomes for *Difference in Arrival Time* (see Figure 8.11) we can see why: its range is more narrow than that of the exploitative policy with a peak at poor values. For the *Time until Containment* objective none of the policies score better than the all-exploitation, though for most the robustness scores lie close (with scores between 0.43-0.48). Only the explorative-185 policy performs notably worse with a score of 0.28. Looking at its distribution, we can see that it has a relatively short tail towards good outcomes compared to the other policies, which explains its poor performance for this robustness measure.

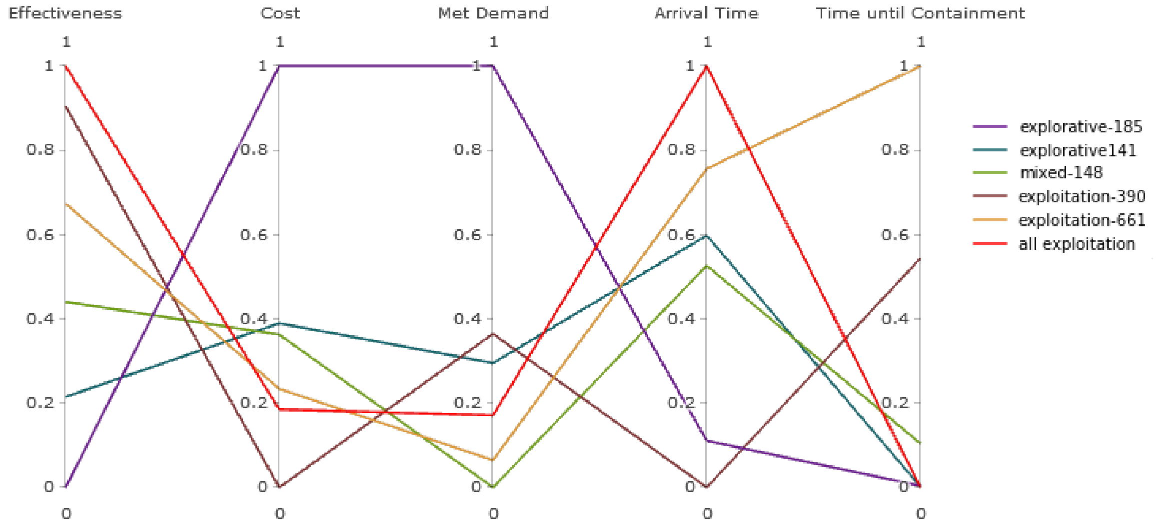

Figure 8.16: Parallel coordinate plot showing the scores of the five policies on the Undesirable Deviations robustness measure. A value of 0 indicates the best score, 1 the worst.

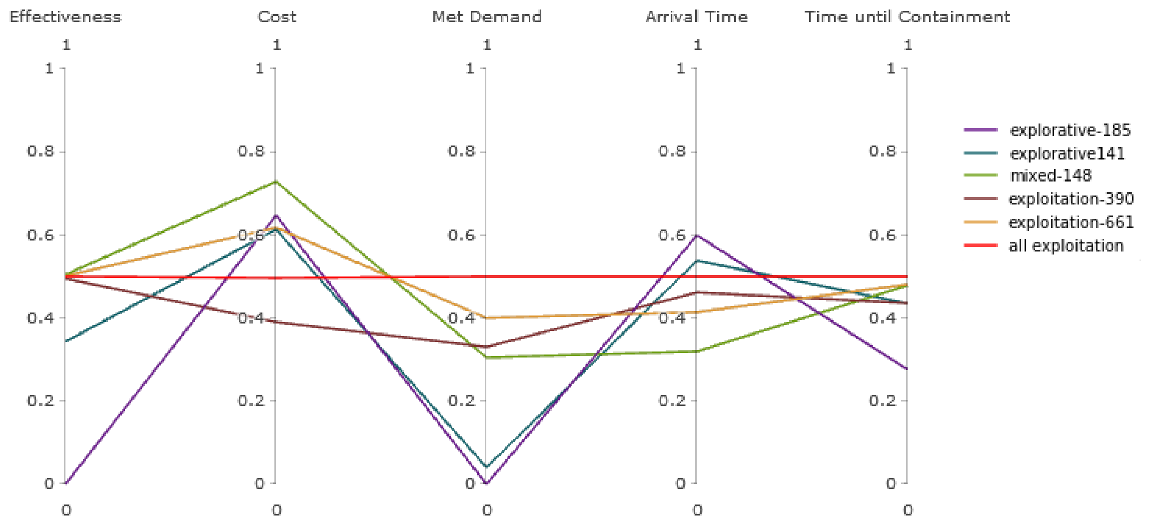

Figure 8.17: Parallel coordinate plot showing the scores of the five policies on the Starr's Domain Criterion robustness measure.

9 Discussion

This chapter provides the interpretation and discussion of the results presented in Chapter 8. The outcomes of the experiments are analysed, placed in to context, and conclusions are drawn from the subsequent discussion. The purpose of this chapter is to formulate answers to the last two sub-questions:

> 5. Given the simulation model, what is the influence of system uncertainties on the performance of resource allocation policies?
>
> 6. Given the simulation model, which strategies for resource allocation decisions show robust performance?

First, the effect of the uncertain variables on the model behaviour and policy performance will be discussed, and related to (epidemiological) conditions which may occur in actual responses. Next, the different types of runtime behaviour generated by the various policies are studied and it is discussed what consequences these have for the performance of a response strategy. The runtime behaviour is also related to the outcome distributions and robustness scores of the different policies. The next section discusses the limitations of the simulation model, both from a technical and humanitarian perspective. Finally, answers to the two sub-questions are formulated in the conclusion.

9.1. Uncertain Factors

One of the first key observations made in Chapter 8 is that the all-exploitation policy's performance is more sensitive to different scenarios than the all-exploration policy is. This is reflected in its outcome distributions, which are much wider, and confirmed by the feature scoring results. To understand why this is the case, we must consider what factors influence how exploitative and explorative decisions are made. In choosing the region in which to place an ETC for an exploitative decision, the decision-maker considers the epidemiological information it has available from all regions. In contrast, when choosing regions which should be explored further, the decision-maker only considers their own experienced level of uncertainty. Naturally, an exploitative decision is therefore more sensitive to factors that impact the dynamics of the epidemic.

For both policies, the transmission rate had the strongest effect on their performance. Low transmission rates were associated with poor performance in terms of effectiveness, and runtime analysis shows why: regardless of the type of response, with low transmission rates the epidemic peaks at the very start of the simulation run, before any ETCs are operational. As a result, a large proportion of the total cases is missed by the response. For the all-exploitation policy, low transmission rates cause another issue: due to the low case numbers, there is a significant chance that no spontaneous news reaches the decision-maker. This results in the epidemic going completely unnoticed.

For the all-exploration policy, the highest range of transmission rate values also lead to poor performance. Because the policy is insensitive towards the environment in which it operates, it always follows the same course of action: until timestep 6, it sends out surveillance teams to explore hidden region. Only after that point in time it will start placing ETCs to further reduce uncertainty, but there is no guarantee that these are placed in the hotspot regions. With high transmission rates, the epidemic will have become so severe at this point in time that it cannot be tackled by the placement of ETCs in random regions.

For the all-exploitation policy however, high transmission rates are associated with an effective response. This is because the explosion of cases will ensure that the decision-maker receives spontaneous news early on in the course of the epidemic, and starts placing treatment centres accordingly. Though the high transmission rate still causes a large number of cases to go untreated by the response (see Figure 8.7), the number of cases the response is able to prevent through early intervention is larger still.

However, these scenarios are also associated with poorer performance in the equity objectives. Even though the decision-maker following an all-exploitation policy has, in these cases, a better situational overview thanks to spontaneous news and an early active presence, scores for equity in met demand and equity in arrival time are low. This is the result of the response remaining concentrated in the hotspot regions, even though the decision-maker is aware of the cases in neighbouring regions. The behaviour seen in the model would be equivalent to a situation in which responders are so overwhelmed by their efforts to meet demand in a certain region, that they cannot aid other areas. However, this can lead to situations where a region with a small number of cases, which could have been easily isolated, is ignored until the region has essentially become a new hotspot. This effect is exacerbated when one of the initial hotspots also has a high number of initial cases.

Given the way in which the response of the all-exploitation policy remains concentrated around the first hotspot regions it discovers, it is unsurprising that a high number of cases in region 4, which is isolated from the other hotspot regions, negatively impacts its ability to contain the epidemic as a whole. In reality, areas of demand could be isolated for a variety of regions - purely on a spatial level, such as in the model, but also because an area has limited telecommunication infrastructure, or because of distrust towards outsiders. Due to its passive nature, a fully exploitative policy often remains unaware of these areas of demand or only stumbles upon them by chance. This can not only lead to the epidemic lasting longer, but also has serious implications in terms of equity, particularly when considering that the reasons that make an area isolated are often related to poverty.

Summarizing, low transmission rates are associated with poor performance for both policy types because the epidemic will have largely burnt out before the response can become properly active. For exploitative policies, there is the additional problem that these types of epidemics can go by completely unnoticed. This is particularly problematic when survivors still experience problems after they have recovered and may require long-term support, as is the case for Ebola. The performance of the fully exploitative policy is also impacted when a significant number of cases exists in an isolated region, or when an epidemic develops quickly. In both cases, due to the passive nature of the response, resources will remain concentrated in and around the first hotspot regions it discovers, leading to poor performance in terms of equity and containment of the epidemic.

9.2. Behaviour and Performance of Policies

In this section, the runtime behaviour of the five policies will be related to their outcome distributions and how the shape of their policy function causes the observed behaviour and outcomes. Based on this discussion some conclusions are drawn on the merits of the different types of policies.

9.2.1. Exploitative Policies

The exploitative-661 policy starts out with a small chance (10%) of taking explorative actions, and in the animated run (see Figure 8.13), this behaviour is visible. Therefore, it is dependent on a combination of explorative actions and the reception of spontaneous news to identify regions in which to place treatment centres. The merit of the explorative actions is based on luck - as soon as some regions have been explored, the policy will become more exploitative, so only a handful of regions can be targeted by these actions. However, as can also be seen in the outcome distribution for *Effectiveness* (Figure 8.11) this does protect the response from scenarios in which the epidemic goes completely unnoticed.

Once the policy becomes fully exploitative, the response remains concentrated on the regions it has already discovered. If isolated regions with a significant number of cases have not been discovered in the earlier stage of the response, they response is likely to remain ignorant of them. As a result, the exploitation-661 policy's performance in terms of equity and containment is very similar to that of the all-exploitation policy.

The exploitation-390 policy (see Appendix H for its runtime behaviour) is almost entirely exploitative at high levels of uncertainty, and is therefore entirely dependent on spontaneous news in order to launch a response. Therefore, this policy also comes with the risk of an epidemic occurring entirely unnoticed, and this is reflected in its distribution of *Effectiveness* scores. Once a response is launched, it remains focussed on placing treatment centres in the regions it is aware of. However, when a tipping point in the level of uncertainty has been reached, it becomes fully explorative almost immediately, and starts "sweeping" the remainder the regions using surveillance teams and later the placement of small treatment centres. This provides the policy with an excellent situational overview later in the response, but since it is now fully explorative, this information cannot be exploited effectively. However, since it now starts placing ETCs in the remainder of the regions systematically, this does mean that the policy performs well in terms of Equity in Arrival Time.

9.2.2. Explorative Policies

As the explorative-185 policy is fully exploitative until uncertainty is reduced to about 75%, it obtains a good situational overview early in the simulation run. However, this knowledge is not exploited in the first half of the simulation run due to is fully-explorative nature. Once its strategy has become mixed, it can target the placement of treatment centres in exploitative decisions well. However, since it is also still performing explorative actions, it is spending a lot of resources in an untargeted manner. These resources are essentially wasted, since explorative actions at this stage of the response do not significantly improve the existing situational awareness of the decision-maker. Additionally, every time it performs an explorative action, it allows the epidemic to develop further in regions which have not received (enough) treatment centres.

The explorative-141 policy also starts out fully-explorative, but moves to a mixed strategy at a much higher level of uncertainty (about 90%). As a result, it also obtains a good situational overview, but is able to start using this information much quicker than the explorative-185 policy. In the simulation runs studied (see Figure 8.15), the explorative-141 policy makes it first exploitative decision to place and ETC at timestep 3, whereas the other explorative policy only does so at timestep 6. This difference in the effective use of situational awareness is also directly visible in the outcome distributions of the two policies. In terms of *Effectiveness*, the explorative-141 policy's optimal performance is significantly better than that of the explorative-185 policy. This difference is even more pronounced in the *Difference in Met Demand* objective. However, the explorative-141 policy does have the same issue as the explorative-185 policy in having a mixed policy for the remainder of the response, therefore wasting resources on explorative actions which do not contribute significantly to the situational awareness.

9.2.3. Mixed Policy

Like the exploitation-661 policy, the mixed-148 policy depends on a combination of explorative actions and the reception of spontaneous news to start its response. However, at this stage it is more explorative than the exploitation-661, and as a result is even better at avoiding the worst outcomes in terms of *Effectiveness*. Yet, because it switches to a fully-exploitative policy at lower levels of uncertainty, it is able to reach almost the same optimal performance of the fully-exploitative policy in terms of *Effectiveness* and *Difference in Met Demand*. Like the exploitative policies, it tends to remain focussed on the hotspot regions and their neighbours as the response evolves. However, thanks to its explorative start it has a high chance of being aware of any isolated regions. As a result, this policy is very active in its placement of treatment centres even later in the response, which is also reflected in its performance in terms of *Time until Containment*, which is better than that of the exploitative policies (i.e. fewer scenarios in which the epidemic is not contained at all).

9.2.4. Value of the Different Policy Types

Overall, it can be concluded that though explorative policies have a much better situational overview throughout the epidemic than the exploitative policies, they fail to use this knowledge to their advantage as they overshoot the balance between explorative and exploitative actions. Likewise, there is no added value in moving to an explorative strategy later in the response to find any remaining pockets of disease, if the incoming information is not used effectively. However, explorative actions have the important property that they protect against the absolute worst case scenario of an epidemic going completely unnoticed.

A mixed approach, with some explorative actions at the start of the response, but an emphasis on exploitative decisions seems to offer sufficient protection against these worst case scenarios while allowing the obtained situational awareness to be used sufficiently. However, both with a fully-exploitative strategy or one that emphasizes exploitation, the response is likely to remain mainly concentrated around the first hotspot regions it becomes aware of. In terms of equity this has a negative effect, particularly if a disease spreads quickly. Though superspreading events have not been fully incorporated in the simulation model (see Chapter 6), it is expected that this focus is more damaging in epidemics in which superspreading events occur regularly. This is the case for any context in which infected individuals are still very mobile (i.e. due to a long incubation period) and in which the infectiousness of the disease is high. In these contexts, the more exploitative policies are “blind” to what happens outside their area of focus, which could have devastating consequences.

9.3. Robustness of the Policies

In the results section, the robustness scores of each of the five policies were presented, with the all-exploitation policy acting as a reference for current practice. These results will be interpreted below, including a discussion of their implications for an actual response.

Undesirable deviations is a regret-based measure - it asks: when it gets bad, how bad does it get? Unsurprisingly, the exploitative policies score poorly in terms of effectiveness. Yet the use of resources for uncertainty reduction, which make the explorative policies robust in terms of effectiveness, make them vulnerable to high costs and poor performance in terms of equity in met demand. This is due to the untargeted way in which these resources are placed. For equity in arrival time, the two policies which are fully exploitative for longer periods of time score best, because they systematically send aid to each region regardless of the epidemiological status of a region. The explorative and mixed policies are also more robust in their ability to contain the epidemic, which can be ascribed to their high level of situational awareness. The all-exploitation policy is the exception here, but is likely caused by an issue in how the *Time until Containment* score is calculated. This is discussed in more detail in the next section.

Overall, the most exploitative (all-exploitation) and explorative (exploration-185) policies are very volatile, scoring well in terms of robustness on some objectives, and performing the worst for others. The mixed policy is, in that sense, the most stable policy in terms of this robustness measure.

Starr’s Domain Criterion is a satisfying measure: it shows if the new policies meet the performance of current practice, and how much better or worse they are. At a first glance, the results are grim: none of the policies meet have higher median values for *Effectiveness*, *Difference in Met Demand*, or *Time until Containment*. For the exploitative and mixed policies however, the difference in the robustness score for *Effectiveness* is very small. Explorative policies never meet the optimal performance in terms of *Effectiveness* and *Difference in Met Demand* because they overshoot the balance of exploring versus exploiting. Because *Difference in Arrival Time* is less sensitive to the proportionality of aid that is received, the explorative policies perform well on this objective because they do establish a presence in many regions relatively quickly. It should also be noted that the scores for *Difference in Met Demand* are somewhat deceptive - the all-exploitation policy, and to a lesser extent the exploitation-661 policy appear to score well on this objective, but these scores are skewed by cases of ‘equity in absence’, which is not a desirable outcome. The exploitation-390 and mixed-148 policies manage to prevent these outcomes. If these equity in absence cases are ignored, the latter policies are likely to score comparatively better in terms of robustness.

Perhaps the most notable outcome for Starr's Domain Criterion is the fact that almost all policies have a lower median in terms of cost. We already saw that the explorative policies had long tails in terms of costs, causing them to score poorly on regret-based robustness in cost. However, Starr's Domain Criterion shows that in favourable scenarios, they are cheaper than the all-exploitation policy. The four policies that are more robust in terms of cost are all policies that start out with non-zero chances of making an explorative decision. Apparently, the increase in situational awareness that this provides means resources can be spent more efficiently later on in the response. Exploitation-390, the only policy to score worse in terms of cost, does the exact opposite: it is fully exploitative early in the response, and fully explorative later. However, these explorative decisions later in the response serve no purpose, and as a result only cause costs to rise.

Starr's domain criterion shows the trade-off between exploitative and explorative decisions well: good situational overview can make a response more efficient and explorative decisions improve equity in arrival times. However, spending resources on these actions in scenarios where the all-exploitative policy performs well means that the optimal performance of that policy becomes unattainable.

In epidemics response, humanitarians and global health actors are likely to be risk-adverse, given the devastating and far-reaching consequences of a failed response. This makes it tempting to ascribe more weight to the Undesirable Deviations robustness scores. However, in reality, the policies that perform best on this measure require resources to be spent on decisions not directly related to providing relief. This can be at odds with the humanitarian principle of humanity, and can also be hard to explain to donors. It is more useful to consider what the two robustness measures represent when both are taken into account: Undesirable Deviations robustness shows what policies are risk-adverse, and the Starr's Domain Criterion robustness shows what the trade-off is for this risk protection in terms of what is lost in optimal performance.

With this in mind, the mixed-148 policy seems to find the balance in averting worst-case scenarios while still maintaining an effective and efficient response. This can attributed to the situational awareness it obtains by making explorative decisions early in the response, before becoming fully exploitative in order to utilize this knowledge effectively.

9.4. Limitations

Several limitations and issues are identified in relation to the simulation model. The first is that the epidemiological model does not make a distinction between rural and urban areas. The fact that higher population densities can cause huge increases in transmissions is therefore not incorporated in the model, though this was a relevant issue in the Ebola response, which saw a concentration of ETCs near cities, a result of the fear of the epidemic becoming uncontrollable in an urban context (Interview E, Interview C). This dimension in the spatial spread of a disease is lost in the model, though it can have a considerable impact on the performance of a policy.

Some artefacts in the results were ascribed to a discrepancy between the definition of the simulation model and the calculation of the objectives. For example, the exclusion of the lowest travel rates for good performance in terms of *Difference in Arrival Time*, as discussed in Chapter 8.1.2. Another issue lies with the formulation of the *Time until Containment* objective: once 70% of the cases are isolated, this objective is considered to be met. However, if the isolation grade falls below this point afterwards, this is not reflected in the model. In the simulation model, this allows for instances in which the objective is met because all cases in region 14 and 15 are isolated early on, while the epidemic could still be active and grow significantly in region 4. Therefore, the number of outcomes in which containment is really achieved is likely smaller than currently represented by the model outcomes, particularly for the explorative policies.

In terms of scope, the present study only considers the isolation and treatment dimension of the response. This is a necessary and reasonable delimitation of the problem given research perspective of uncertainty reduction as a result of resource allocation as discussed in Chapter 3 and 6. However, the focus of this scope limits the extent to which the results can be translated to a real-life epidemic, particularly in the case of Ebola. Fear and distrust played a huge role in the spread of the epidemic in 2014 (see for example Interview B), and these same issues are arising again in the current outbreak in the DRC (NOS, 2019). Community engangement was seen as a key component in stopping the disease (M'Cormack-Hale et al., 2016; Médecins Sans Frontières, 2015). Perhaps most telling is that almost all interview subjects, who had all been involved in the 2014 Ebola epidemic, brought up the relevance of social and cultural factors in epidemics response without prompting. For example, in interview A, the interviewee indicated *"Now it seems like lots and lots of resources were assigned to treat people, rather than preventing people from getting Ebola [...] they were looking for data in order to set up ETUs and so on, but maybe they should instead have done more education activities"*. Another expert cited the point at which traditional burial rituals could be curbed as a turning point, saying the extent to which community engagement played a role was a big eye-opener for them (Interview B).

The fact that these components are not incorporated in the model does not invalidate the model outcomes on a theoretical level. It does, however, mean that strategies identified as a result of the model can likely never be implemented optimally. In the words of one of the experts while the simulation model was being discussed: *"You can make models with all the possible parameters, but if a community says no, then it's really difficult"* (Interview C). This is a major concern for the translation between research and practice, since analysis has shown that small differences in policies can cause large differences in outcomes (consider for example the big difference in performance between policy 185 and policy 141 in terms of *Difference in Met Demand*). Ensuring the limitations and assumptions associated with the research results are communicated well to practitioners is therefore an important aspect in bridging the gap between theory and practice (see also Interview F).

Another relevant limitation in terms of bridging the gap between research and practice is that some policies might not be implemented on basis of principle. It is unlikely that humanitarian organisations would be willing to participate in a fully explorative policy if it means knowingly ignoring identified demand. Yet several of the identified policies require a fully explorative strategies at different moments in time. A situation which might be acceptable on the basis of principle, is taking an explorative approach late in the response, when aid has already been provided to several regions. However, this has been shown to be highly inefficient and resource intensive, which might make it hard to sell to national governments, WHO Member States and humanitarian donors alike. These issues need to be carefully considered when attempting to translate the research results to a more practical application.

9.5. Conclusion

The analysis of the results has provided the necessary information to answer the two sub-questions:

> 5. Given the simulation model, what is the influence of system uncertainties on the performance of resource allocation policies?

The way in which the uncertain factors influence policy performance is dependent on the type of policy. For both exploitative and explorative policies, the transmission rate is the most influential factor. For the all-exploitation policy, high transmission rates are associated with good performance in terms of effectiveness and efficiency, since the explosive nature of the epidemic means it will be discovered quickly. However, high transmission rates have a negative effect on equity as these will cause the response to remain concentrated on the hotspot regions. Since the all-exploration policy does not use the information it obtains through its uncertainty-reducing actions, high transmission rates cause poor performance since the epidemic will grow beyond control. The lowest transmission rates are associated with poor performance in terms of effectiveness and efficiency for both policies, as the epidemic will have already peaked by the time the first ETCs become operational. For the all-exploitation policy, these conditions also bring the risk of the epidemic going completely unnoticed. Because the all-exploitation policy often does not discover isolated hotspot regions, a high number of initial cases in such a region negatively affects its ability to contain the epidemic quickly. Likewise, a high number of cases in hotspot regions can cause the all-exploitation policy to perform poorly in terms of equity, as that will cause efforts to be concentrated on those regions. These factors do not affect the all-exploration policy, as the policy's resource allocation strategy does not use epidemiological information to determine where to send resources. This also explains why the all-exploration policy's performance is less sensitive to the uncertain factors than the all-exploitation policy.

> 6. Given the simulation model, which strategies for resource allocation decisions show robust performance?

The robustness of exploitative, mixed, and explorative strategies differs depending on what robustness metrics are used. Generally, explorative policies are better at preventing the worst outcomes, but this comes at the cost of lower optimal performance in the best scenarios due to the trade-off of spending resources on exploration as well as direct treatment. Explorative policies are more robust in their performance in terms of equity in arrival times, regardless of which measure is used. Overall, mixed policies with an emphasis on exploitation at lower levels of uncertainty seem to be most robust. Their explorative element prevents an epidemic from being discovered too late or not at all, while the shift towards exploitative decisions later on in the response ensures known cases are treated.

10 Conclusion

10.1. Summary

This thesis has developed a spatial resource allocation model combined with a dynamic epidemiological model in order to simulate sequential resource allocation decisions during an epidemic. Within this model, deep uncertainty was incorporated in two ways: first, several uncertain input variables were identified based on the case study of the 2014 Ebola epidemic. The effect of variation in these uncertain factors was studied by considering multiple objectives related to the performance of allocation strategies. Additionally, uncertainty was incorporated in-model by including uncertainty in the information the decision-maker has available to make allocation decisions during runtime. Experienced uncertainty can be reduced as the result of decisions.

A conceptualization of this uncertainty reduction was proposed, where a uniform range around a ground truth is decreased based on factors relevant to the uncertain variable, i.e. time, the number of cases, or the presence of an ETC in a neighbouring region. This conceptualization was translated into the simulation model, in which each time step either an explorative action aimed at reducing uncertainty can be taken, or an exploitative action aimed at isolating and treating cases. Which type of action is taken is determined by a policy function which is dependent on the level of uncertainty experienced by the decision-maker. By doing so it adds another dimension of deep uncertainty to the method of direct policy search.

Exploratory modelling analysis has shown that the transmission rate is the most influential factor in determining policy performance, as it is the dominant factor in terms of epidemiological dynamics. High transmission rates are associated with good performance by exploitative policies, as “explosive” epidemics have a higher chance of being discovered. Due to its passive nature, the exploitative policy is also sensitive to the initial distribution of cases over the regions. A high number of cases in an isolated region causes the policy to fail to contain the epidemic, and if the initial situation in one of the hotspot regions is severe, equity is lower as the response remains focussed on those regions. Since it does not base its placement decisions on epidemiological factors, the all-exploration policy is less sensitive to the uncertain factors. However, the transmission rate still has a considerable effect, since it determines whether the epidemic can be contained with the untargeted placement of treatment centres later in the response.

The different policy functions found by the MOEA used for policy search confirm that the connection between experienced uncertainty and decision strategies is relevant for epidemiological resource allocation. Explorative policies tend to be more robust in terms of risk aversion and lead to greater situational awareness, but their optimal performance is not as high as that of the exploitative policies. This is because they operate on a trade-off between using resources for isolation and treatment, and using them for exploration. Explorative policies show better performance in equity in arrival times because they distribute resources more evenly over the regions. However, despite having a better situational overview on the location and magnitude of demand than exploitative policies, explorative policies do not perform better in terms of equity in met demand because they spend too many resources on explorative decisions rather than meeting known demand.

10.2. The value of explicitly incorporating uncertainty in resource allocation strategies

In the previous chapters, each of the sub-questions formulated in Chapter 2 were answered. This provides the information necessary to answer the main research question.

> What is the value of explicitly incorporating uncertainty and its reduction by sequential treatment centre placement decisions for epidemics response?

Primarily, the value of considering uncertainty and aiming to reduce it by taking explorative actions is in reducing the risk of an epidemic being discovered too late, or not at all. This is most relevant when epidemics develop slowly or when the regions in which the epidemic is taking place are isolated from one another. However, since there is always a trade-off between how resources are spent, taking explorative actions in a quickly evolving epidemic may actually lead to poorer performance, since more lives could be saved by fully targeting the response on known cases.

Explorative strategies are also associated with higher equity in arrival times, because they distribute resources over more regions. They also generate a better situational overview. Therefore, aiming to reduce uncertainty is also valuable when considering the ethical principles associated with humanitarianism, because they provide decision-makers with a better overview of the situation to which they are responding, allowing them to weigh different response options more carefully. Here the trade-off between exploring and exploiting can also be problematic, as was seen for the equity in met demand objective. Any benefit of having a better situational overview due to explorative actions is negated because resources are not fully spent on meeting demand, causing the situation in hotspot regions to spiral out of control while other regions are being explored. However, exploitative strategies can also prove to be problematic for this objective, as they perform particularly poorly in contexts where some regions are more likely to draw attention than others in the event of an epidemic, which could be the case for richer regions, cities, or more densely populated areas.

When combining explorative and exploitative strategies in a response, the sequence in which the types of decisions are made is important. In any case, making explorative decisions is only valuable if they are, at some point in the future, followed by exploitative decisions. Otherwise, the (improved) situational awareness gained from uncertainty reduction is never used for targeted placement decisions. The policies studied in this thesis suggest that a mixed strategy early in the response provides protection against the worst case scenarios, but the remainder of the response should be fully exploitative to ensure the discovered cases are effectively treated.

10.3. Scientific Contribution

The scientific contribution of this thesis is two-fold. Its first contribution is that it applies deep uncertainty methods to dynamic resource allocation for epidemics. As discussed in the literature review in the first chapter, the effect different parametrizations of a model can have on the performance of an allocation strategy are rarely considered systematically in resource allocation research for epidemics. The outcomes of this thesis show that it is relevant to include them, because the performance of the all-exploitation policy depended on the transmission rate and initial distribution of cases. Additionally, the model developed in this thesis included multiple objectives, providing more in-depth insight in benefits and drawbacks of different policies and an indication to which objectives are strongly related.

The second contribution of this thesis relates to its incorporation of uncertainty reduction as a result of allocation decisions. This interaction was first considered for humanitarian logistics in a master thesis by Romijn (2018), and is expanded on in this thesis by using the level of uncertainty as the variable on which policy functions are based. By creating this “in-model” feedback mechanism, it introduces a new dimension to the closed loop policies designed with direct policy search in deep uncertainty research.

As this method provided valuable insights, proving its validity, it has opened the door to a new way of handling uncertainty within deep uncertainty research.

10.4. Societal Relevance

By exploring the strengths and vulnerabilities of an all-exploitation policy, which can be considered as an approximation of current practice, this thesis has shown how disease and context-specific factors relate to the performance of such an allocation strategy. This provides valuable insights to global health actors and humanitarians alike, by showing in which types of scenarios current practice is at risk of performing poorly in terms of effectiveness or equity. Additionally, by considering different strategies for resource allocation, which also incorporate explorative actions to obtain more information and reduce uncertainty, this thesis provides the first insights in which strategies can be used to address the problems associated with exploitative policies. Most significantly, incorporating explorative actions can prevent worst-case scenarios in which resource allocation strategies fail to provide treatment, causing an epidemic to spiral out of control. In light of the responsibility that the international community and humanitarian organisations bear towards those in need, this is an important finding.

10.5. Future Research

Direct continuations of this work would be to study how the mixed policies are influenced differently by the model uncertainties, which would lead to a deeper understanding of their performance under varying conditions. Another model component that is worthy of more attention is the spontaneous news mechanism. As outlined in the discussion, its impact on model outcomes is significant, which makes investigating model sensitivity to its function worthwhile. This could include different parametrisations of the sigmoidal function which is currently used to represent this mechanism, as well as an investigation into different function shapes or even completely different mechanisms. Additionally, the policy function could be developed further. Currently, it is only dependent on the observed level of uncertainty, which means that sometimes it makes choices which are obviously bad. Creating a two-dimensional policy function which also takes into account the perceived number of cases is likely to improve the quality of the policies significantly.

To my best knowledge, this study is the first to incorporate deep uncertainty in a spatial resource allocation problem using a dynamic epidemiological model. The first steps made in this study have shown the value of this approach, while throughout the research process many different aspects that need further attention have been identified. Three perspectives for future research are outlined here.

The first avenue of future research is within the paradigm of deep uncertainty. This study showed that explicitly incorporating system uncertainty from the perspective of the decision-maker in an adaptive policy is valuable. Since using the level of uncertainty as an input factor for policies is a novel idea, there are many aspects which still require further research. One is to study how uncertainty is experienced by the decision-maker and how uncertainty is incorporated in their decisions. For example, in this thesis the decision-maker had an estimated range with a uniform probability available for relevant variables for decision-making, and the lower bound of this distribution was used to predict the number of future cases. However, many different conceptualizations for the perception and use of uncertainty by the decision-maker are possible. These can be context-dependent, though it is also worthwhile to investigate how different conceptualizations for the same problem affect performance. Another aspect worth of future study is the reduction of uncertainty as a result of decisions. Clearly, if policies are dependent on the level of uncertainty, uncertainty is expected to change over the course of the decision-making window. However, how this uncertainty changes exactly needs to be studied further. This thesis made a first attempt to base uncertainty reduction on factors supported by evidence, but there is a lot of room for further improvement of this method. Additionally, in a dynamic environment it is also possible for uncertainty to increase again after it is reduced. This adds a new dimension to the interaction between decisions and the environment. Further study into any of these elements would improve understanding on the interaction between decisions and changes in uncertainty. Finally, the approach developed in this thesis can potentially be applied to many different decision problems, given that there is an interaction between the level of uncertainty in the system and decisions made. Applying the approach in different areas of decision making will create an understanding on which parts of the method can be generalized, and which parts require problem-specific tailoring. This can improve the effectiveness of the approach and can also provide insights into the type of contexts or problems for which it is valuable.

The second direction is in epidemiological modelling and response. This thesis has shown how a response is affected by uncertain factors and how these relate to (un)desirable outcomes. The uncertain factors used for this were chosen on the basis of the case study of the Ebola crisis, and might not be present or relevant in other epidemics. Therefore, research aiming at identifying uncertain factors related to different (types of) epidemics, and applying exploratory modelling techniques to study their effect can lead to a better generalized understanding of how epidemics can evolve and what this means for response performance.

Finally, from the perspective of humanitarian decision-makers the implications of exploitative versus explorative strategies can be investigated further. The trade-off between these strategies as identified in this research raises value-based questions about responsibility toward affected populations in combination with the decision-maker's awareness of their suffering. To what extent are humanitarian decision-makers responsible not just for providing relief to those in need, but also for finding affected populations? Additionally, the difference in robustness for exploitative and explorative strategies also raises questions about risk, and how much risk is acceptable when pursuing the best results. Understanding the implications of these different policies is a vital step in bridging the gap between theory and practice.

References

Adams, V., Allié, M.-P., Burrier, P., Biedler, M., Bonis, E., de, Colebunder, R., … Sánchez, V. (2016, April). *OCB Ebola Review: Summary Report* (Tech. Rep.). MSF - Stockholm Evaluation Unit. Retrieved from https://evaluation.msf.org/sites/evaluation/files/attachments/ocb_ebola_review_summary_report_final_3.pdf

Ahmadi-Javid, A., Seyedi, P., & Syam, S. S. (2017). A survey of healthcare facility location. *Computers & Operations Research*, *79*, 223–263.

Alirol, E., Getaz, L., Stoll, B., Chappuis, F., & Loutan, L. (2011). Urbanisation and infectious diseases in a globalised world. *Lancet Infect Dis*, *11*(2), 131-41. doi: 10.1016/s1473-3099(10)70223-1

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

Atkins, K., E, Pandey, A., Wenzel, N., S, Skrip, L., Yamin, D., Nyenswah, T., G, … Altice, F., L (2016). Retrospective analysis of the 2014–2015 Ebola epidemic in Liberia. *The American journal of tropical medicine and hygiene*, *94*(4), 833-839.

Auping, W. L., Pruyt, E., & Kwakkel, J. H. (2017). Simulating endogenous dynamics of intervention-capacity deployment: Ebola outbreak in liberia. *International Journal of Systems Science: Operations & Logistics*, *4*(1), 53–67.

Backer, J. A., & Wallinga, J. (2016, 12). Spatiotemporal analysis of the 2014 ebola epidemic in west africa. *PLOS Computational Biology*, *12*(12), 1-17. Retrieved from https://doi.org/10.1371/journal.pcbi.1005210 doi: 10.1371/journal.pcbi.1005210

Balcik, B., & Beamon, B. M. (2008). Facility location in humanitarian relief. *International Journal of Logistics*, *11*(2), 101–121.

Bevere, L., Schwartz, M., Sharan, R., & Zimmerli, P. (2018). Natural catastrophes and man-made disasters in 2017: a year of record-breaking losses. *Sigma*(1). Retrieved from https://reliefweb.int/sites/reliefweb.int/files/resources/sigma1_2018_en.pdf

Bloom, D. E., Cadarette, D., & Sevilla, J. . (2018, June). Epidemics and Economics. *Finance & Development*, *55*(2). Retrieved from https://www.imf.org/external/pubs/ft/fandd/2018/06/economic-risks-and-impacts-of-epidemics/bloom.htm

Brauer, F. (2008). Compartmental models in epidemiology. In F. Brauer, P. van den Driessche, & J. Wu (Eds.), *Mathematical epidemiology* (pp. 19–79). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from https://doi.org/10.1007/978-3-540-78911-6_2 doi: 10.1007/978-3-540-78911-6_2

Bryant, B. P., & Lempert, R. J. (2010). Thinking inside the box: a participatory, computer-assisted approach to scenario discovery. *Technological Forecasting and Social Change*, *77*(1), 34–49.

Büyüktahtakın, . 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*(3), 1046 - 1063. Retrieved from http://www.sciencedirect.com/science/article/pii/S0377221717307683 doi: https://doi.org/10.1016/j.ejor.2017.08.037

Çelik, M., Ergun, Ö., Johnson, B., Keskinocak, P., Lorca, Á., Pekgün, P., & Swann, J. (2012). Humanitarian logistics. In *New directions in informatics, optimization, logistics, and production* (pp. 18–49). INFORMS.

Centers for Disease Control and Prevention. (2011). *Cholera in Haiti: One Year Later* (Vol. 2019) (Web Page No. 29 January). Website. Retrieved from https://www.cdc.gov/cholera/haiti/haiti-one-year-later.html

Centers for Disease Control and Prevention. (2019, January). *Ebola (Ebola Virus Disease) - Prevention.* Retrieved from https://www.cdc.gov/vhf/ebola/prevention/index.html (Accessed on 15 April 2019)

Charles, A., & Lauras, M. (2011). An enterprise modelling approach for better optimisation modelling: application to the humanitarian relief chain coordination problem. *OR spectrum*, *33*(3), 815–841.

Dalal, S., Han, B., Lempert, R., Jaycocks, A., & Hackbarth, A. (2013). Improving scenario discovery using orthogonal rotations. *Environmental modelling & software*, *48*, 49–64.

Dasaklis, T. K., Pappis, C. P., & Rachaniotis, N. P. (2012). Epidemics control and logistics operations: A review. *International Journal of Production Economics*, *139*(2), 393–410.

Dasaklis, T. K., Rachaniotis, N., & Pappis, C. (2017). Emergency supply chain management for controlling a smallpox outbreak: the case for regional mass vaccination. *International Journal of Systems Science: Operations & Logistics*, *4*(1), 27–40.

Diallo, B., & Dilorenzo, S. (2014, September 8 2014). *Ebola virus is surging in places where it was beaten back: experts* (Vol. 2019) (Web Page No. 18 Febuary). Retrieved from `https://www.ctvnews.ca/health/ebola-virus-is-surging-in-places-where-it-was-beaten-back-experts-1.1996179`

Gettleman, J. (2014, Dec. 7). *As ebola rages, poor planning thwarts efforts* [Newspaper Article]. Retrieved from `https://www.nytimes.com/2014/12/07/world/africa/as-ebola-rages-in-sierra-leone-poor-planning-thwarts-efforts.html?action=click&module=RelatedCoverage&pgtype=Article&region=Footer`

Gholipour, B. (2013, 26 November). *What 11 Billion People Mean for Disease Outbreaks* [Newspaper Article]. Retrieved from `https://www.scientificamerican.com/article/what-11-billion-people-mean-disease-outbreaks/`

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P., M. (2015). Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations. *Journal of Water Resources Planning and Management*, *142*(2).

Haasnoot, M., Warren, A., & Kwakkel, J. H. (2019). Dynamic adaptive policy pathways (dapp). In V. A. W. J. Marchau, W. E. Walker, P. J. T. M. Bloemen, & S. W. Popper (Eds.), *Decision making under deep uncertainty: From theory to practice* (pp. 71–92). Cham: Springer International Publishing. Retrieved from `https://doi.org/10.1007/978-3-030-05252-2_4` doi: 10.1007/978-3-030-05252-2_4

Hadka, D., & Reed, P. (2013). Borg: An auto-adaptive many-objective evolutionary computing framework. *Evolutionary computation*, *21*(2), 231–259.

Harman, S., & Wenham, C. (2018). Governing ebola: between global health and medical humanitarianism. *Globalizations*, *15*(3), 362-376. Retrieved from `https://doi.org/10.1080/14747731.2017.1414410` doi: 10.1080/14747731.2017.1414410

He, Y., & Liu, N. (2015). Methodology of emergency medical logistics for public health emergencies. *Transportation Research Part E: Logistics and Transportation Review*, *79*, 178–200.

Hilhorst, D. (2005). Dead letter or living document? Ten years of the Code of Conduct for disaster relief. *Disasters*, *29*(4), 351–369. Retrieved 2019-03-20TZ, from `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0361-3666.2005.00297.x` doi: 10.1111/j.0361-3666.2005.00297.x

Holguín-Veras, J., Jaller, M., Van Wassenhove, L. N., Pérez, N., & Wachtendorf, T. (2012). On the unique features of post-disaster humanitarian logistics. *Journal of Operations Management*, *30*(7-8), 494–506.

Huang, K., Jiang, Y., Yuan, Y., & Zhao, L. (2015, March). Modeling multiple humanitarian objectives in emergency response to large-scale disasters. *Transportation Research Part E: Logistics and Transportation Review*, *75*, 1–17. Retrieved 2019-03-24TZ, from `http://www.sciencedirect.com/science/article/pii/S136655451400204X` doi: 10.1016/j.tre.2014.11.007

Kaner, J., & Schaack, S. (2016). Understanding Ebola: The 2014 Epidemic. *Globalization and health*, *12*(1), 53.

Karlsen, A., & Kruke, B. I. (2018). The role of uncertainty during the ebola pandemic in western africa (2014–2016). *Journal of Extreme Events*, *05*(01), 1850009. Retrieved from `https://doi.org/10.1142/S2345737618500094` doi: 10.1142/S2345737618500094

Kucharski, A. J., Camacho, A., Flasche, S., Glover, R. E., Edmunds, J. W., & Funk, S. (2015). Measuring the impact of Ebola control measures in Sierra Leone. *Proceedings of the National Academy of Sciences*, *112*(46), 14366-14371.

Kwakkel, J. H. (2017). The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making. *Environmental modelling & software*, *96*, 239–250.

Lempert, R. J. (2003). *Shaping the next one hundred years: new methods for quantitative, long-term policy analysis*. Rand Corporation.

Liberatore, F., Pizarro, C., de Blas, C. S., Ortuño, M. T., & Vitoriano, B. (2013). Uncertainty in Humanitarian Logistics for Disaster Management. A Review. In B. Vitoriano, J. Montero, & D. Ruan (Eds.), *Decision aid models for disaster management and emergencies* (pp. 45–74). Paris: Atlantis Press. Retrieved from https://doi.org/10.2991/978-94-91216-74-9_3 doi: 10.2991/978-94-91216-74-9_3

Liu, M., & Zhang, D. (2016). A dynamic logistics model for medical resources allocation in an epidemic control with demand forecast updating. *Journal of the Operational Research Society*, *67*(6), 841–852.

Mbah, M. L. N., & Gilligan, C. A. (2011). Resource allocation for epidemic control in metapopulations. *PLoS One*, *6*(9), e24577.

McNiel, D. G., Jr. (2015, January). *Ebola Doctors Are Divided on IV Therapy in Africa.* Online Newspaper Article. Retrieved from https://www.nytimes.com/2015/01/02/health/ebola-doctors-are-divided-on-iv-therapy-in-africa.html (Accessed 19 April 2019)

M'Cormack-Hale, F., Lavali, A., & Magbity, E. (2016, March). *Dutch Relief Assistance to Ebola in West Africa: An Independent Evaluation* (Tech. Rep.). Institute for Governance Reform (IGR), Sierra Leone. Retrieved from https://www.terredeshommes.nl/sites/tdh/files/visual_select_file/copy_of_nl_2017_05_15_jre_final_evaluation_report.pdf

McPhail, C., Maier, H., Kwakkel, J., Giuliani, M., Castelletti, A., & Westra, S. (2018). Robustness metrics: How are they calculated, when should they be used and why do they give different results? *Earth's Future*, *6*(2), 169–191.

Médecins Sans Frontières. (2015). *Ebola: Pushed to the limit and beyond.* Retrieved from https://www.msf.org/ebola-pushed-limit-and-beyond (Accessed April 5 2019)

Meltzer, M. I., Atkins, C. Y., Santibanez, S., Knust, B., Petersen, B. W., Ervin, E. D., … Washington, M. L. (2014). Estimating the future number of cases in the Ebola epidemic–Liberia and Sierra Leone, 2014–2015. *MMWR. Morbidity and Mortality Weekly Report*, *63*(3).

Memarzadeh, M., & Pozzi, M. (2016). Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling. *Reliability Engineering & System Safety*, *154*, 137-151.

Moon, S., Sridhar, D., Pate, M. A., Jha, A. K., Clinton, C., Delaunay, S., … Piot, P. (2015, November). Will Ebola change the game? Ten essential reforms before the next pandemic. The report of the Harvard-LSHTM Independent Panel on the Global Response to Ebola. *The Lancet*, *386*(10009), 2204–2221. Retrieved CURRENT_TIMESTAMP, from https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(15)00946-0/abstract doi: 10.1016/S0140-6736(15)00946-0

Murali, P., Ordóñez, F., & Dessouky, M. M. (2012). Facility location under demand uncertainty: Response to a large-scale bio-terror attack. *Socio-Economic Planning Sciences*, *46*(1), 78–87.

NOS. (2019). *Ebola in Congo: waarom is het wantrouwen zo groot? [Ebola in Congo: why is distrust so severe?].* Online News Article. Retrieved from https://nos.nl/artikel/2288916-ebola-in-congo-waarom-is-het-wantrouwen-zo-groot.html

Nyenswah, T., Fahnbulleh, M., Massaquoi, M., Nagbe, T., Bawo, L., Falla, J. D., … others (2014). Ebola epidemic—Liberia, March–October 2014. *MMWR. Morbidity and Mortality Weekly Report*, *63*(46), 1082.

Payne, E. (2014, September 12 2014). *West African health centres can't keep up with Ebola oubeak, WHO says* (Vol. 2019) (Web Page No. 19 Febuary). Retrieved from https://edition.cnn.com/2014/09/12/health/ebola-outbreak/index.html

Pianosi, F., Beven, K., Freer, J., Hall, J. W., Rougier, J., Stephenson, D. B., & Wagener, T. (2016). Sensitivity analysis of environmental models: A systematic review with practical workflow. *Environmental Modelling & Software*, *79*, 214–232.

Quinn, J., D, Reed, P., M, & Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environmental Modelling & Software*, *92*, 125-141.

Rachaniotis, N. P., Dasaklis, T. K., & Pappis, C. (2017). Controlling infectious disease outbreaks: A deterministic allocation-scheduling model with multiple discrete resources. *Journal of Systems Science and Systems Engineering*, *26*(2), 219–239.

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*(1), 225–231.

Ren, Y., Ordóñez, F., & Wu, S. (2013). Optimal resource allocation response to a smallpox outbreak. *Computers & Industrial Engineering*, *66*(2), 325–337.

Rivers, C. M., Lofgren, E. T., Marathe, M., Eubank, S., & Lewis, B. L. (2014). Modeling the impact of interventions on an epidemic of ebola in sierra leone and liberia. *PLoS currents*, *6*.

Romijn, T. (2018). *Deep Uncertainty in Humanitarian Logistics: Simulation and Analysis of the Interplay between Decisions and Uncertainty for Post-Disaster Facility Location Decisions* (Unpublished master's thesis). Delft University of Technology.

Scarpino, S. V., Iamarino, A., Wells, C., Yamin, D., Ndeffo-Mbah, M., Wenzel, N. S., … others (2014). Epidemiological and viral genomic sequence analysis of the 2014 Ebola outbreak reveals clustered transmission. *Clinical Infectious Diseases*, *60*(7), 1079–1082.

Secretariat of the World Health Organization. (2015). *The role of who in the united nations mission for ebola emergency response*. WHO. Retrieved from `https://www.who.int/csr/resources/publications/ebola/who-unmeer.pdf` (Accessed on 20 May 2019)

Sánchez Carrera, V. (2015). *EBOLA REVIEW Ebola Treatment Centres: design and construction PART I –Evaluation outcomes* (Tech. Rep.). MSF. Retrieved from `http://evaluation.msf.org/sites/evaluation/files/attachments/seu_ebola_report_etc_design_and_construction_part_i_final_0.pdf`

Sphere Association, et al. (2018). *Sphere handbook: Humanitarian charter and minimum standards in humanitarian response*. PRACTICAL ACTION.

Statistics Sierra Leone. (n.d.). *2015 population and housing census summary of final results* [Online PDF].

United Nations. (2017, 21 June 2017). *World population projected to reach 9.8 billion in 2050, and 11.2 billion in 2100* [Press Release]. Webpage. Retrieved from `https://www.un.org/development/desa/en/news/population/world-population-prospects-2017.html`

United Nations. (2019). *UN Mission for Ebola Emergency Response (UNMEER).* Retrieved from `https://ebolaresponse.un.org/un-mission-ebola-emergency-response-unmeer` (Accessed on 20 March 2019)

UNOCHA. (2014, September). *Sierra leone: Humanitarian operational presence against the ebola epidemic (as of 1 sep 2014).* Retrieved from `https://reliefweb.int/map/sierra-leone/sierra-leone-humanitarian-operational-presence-against-ebola-epidemic-3-sep-2014` (Accessed on 17 June 2019)

van der Laan, E., van Dalen, J., Rohrmoser, M., & Simpson, R. (2016). Demand forecasting and order planning for humanitarian logistics: An empirical assessment. *Journal of Operations Management*, *45*, 114–122.

Vetter, P., Dayer, J.-A., Schibler, M., Allegranzi, B., Brown, D., Calmy, A., … Pittet, D. (2016, May 05). The 2014–2015 ebola outbreak in west africa: Hands on. *Antimicrobial Resistance & Infection Control*, *5*(1), 17. Retrieved from `https://doi.org/10.1186/s13756-016-0112-9` doi: 10.1186/s13756-016-0112-9

Wanying, C., Alain, G., & Angel, R. (2016). Modeling the logistics response to a bioterrorist anthrax attack. *European Journal of Operational Research*, *254*(2), 458–471.

WHO Ebola Response Team. (2014). Ebola virus disease in West Africa—the first 9 months of the epidemic and forward projections. *New England Journal of Medicine*, *371*(16), 1481–1495.

WHO Ebola Response Team. (2016). After ebola in west africa — unpredictable risks, preventable epidemics. *New England Journal of Medicine*, *375*(6), 587-596. Retrieved from `https://doi.org/10.1056/NEJMsr1513109` (PMID: 27509108) doi: 10.1056/NEJMsr1513109

World Health Organization. (2003). *Climate change and human health* (Report). Author. Retrieved from `http://www.who.int/globalchange/climate/summary/en/index5.html`

World Health Organization. (2014a). *Basic documents* (48th ed.). WHO,. Retrieved from `http://apps.who.int/gb/bd/PDF/bd48/basic-documents-48th-edition-en.pdf#page=1` (Accessed on April 5 2019)

World Health Organization. (2014b). *Ebola Responde Roadmap (August 28 2014)*. World Health Organization. Retrieved from `https://apps.who.int/iris/bitstream/handle/10665/131596/EbolaResponseRoadmap.pdf?sequence=1` (Accessed April 5 2019)

World Health Organization. (2014c, 10 December). *Ebola Response Roadmap Situation*

Report (December 10) (Report). Author. Retrieved from https://apps.who.int/iris/bitstream/handle/10665/145198/roadmapsitrep_10Dec2014_eng.pdf;jsessionid=6C37FB855BD0CD36412097F3BE49302B?sequence=1

World Health Organization. (2014, August). *Ebola Strategy - Ebola and Marburg virus disease epidemics: preparedness, alert, control and evaluation* (techreport). World Health Organization,. Retrieved from https://apps.who.int/iris/bitstream/handle/10665/130160/WHO_HSE_PED_CED_2014.05_eng.pdf?sequence=1 (Accessed 10 April 2019)

World Health Organization. (2014, September). *WHO: Ebola response roadmap update 26 september 2014.* Retrieved from https://apps.who.int/iris/bitstream/handle/10665/135029/roadmapupdate26sept14_eng.pdf (Accessed on 21 March 2019)

World Health Organization. (2016). *Ebola data and statistics - Latest Sierra Leone data in CSV, as of May 11 2016.* CSV-file. Retrieved from http://apps.who.int/gho/data/view.ebola-sitrep.ebola-country-GIN-latest-download

World Health Organization. (2018, October). *Ebola virus disease - Frequently Asked Questions.* Retrieved from https://www.who.int/ebola/drc-2018/faq-vaccine/en/ (Accessed on 15 April 2019)

World Health Organization. (2019). *What We Do.* Web page. Retrieved from https://www.who.int/about/what-we-do (Accessed April 5 2019)

World Health Organization, & Centers for Disease Control and Prevention. (2015). *Implementation and management of contact tracing for Ebola virus disease* (Tech. Rep.). World Health Organization. Retrieved from https://apps.who.int/iris/bitstream/handle/10665/185258/WHO_EVD_Guidance_Contact_15.1_eng.pdf?sequence=1 (Accessed 10 April 2019)

World Health Organization and others. (2014). *Ebola and Marburg virus disease epidemics: preparedness, alert, control, and evaluation* (Tech. Rep.). World Health Organization. Retrieved from https://apps.who.int/iris/bitstream/handle/10665/130160/WHO_HSE_PED_CED_2014.05_eng.pdf?sequence=1 (Accessed 10 April 2019)

Yang, J., & Leskovec, J. (2010, December). Modeling Information Diffusion in Implicit Networks. In *2010 IEEE International Conference on Data Mining* (p. 599-608). doi: {10.1109/ICDM.2010.22}

A

List of Abbreviations used in this Thesis

CDC Centre for Disease Control

CI Confidence Interval

ETC Ebola Treatment Centre (equivalent to ETU, an Ebola Treatment Unit)

EVD Ebola Virus Disease

DPS Direct Policy Search

IFRC International Federation of Red Cross and Red Crescent Societies

MOEA Many-Objective Evolutionary Algorithm

MORDM Many-Objective Robust Decision-Making

NGO Non-Governmental Organisation

MSF Médecines sans Frontières

PCA Principal Component Analysis

PRIM Patient Rule Induction Method

UN United Nations

UNMEER United Nations Mission for Ebola Emergency Response

UNOCHA United Nations Office for the Coordination of Humanitarian Affairs

WHO World Health Organization

B

Uncertainty Reduction - Lookup Tables and Functions used in the Simulation Model

Lookup Tables for the reduction in uncertainty on the number of infections in a region:

Uncertainty reduction for a 100-bed ETC:

Timsteps after placement decision	1	2	3	4	5	6	7	8	9	10	11	12
% of uncertainty remaining	100	95	90	85	75	70	50	32.5	25	22.5	21	20

Uncertainty reduction for 50-bed ETC:

Timsteps after placement decision	1	2	3	4	5	6	7	8	9	10
% of uncertainty remaining	100	95	85	70	50	32.5	25	22.5	21	20

Uncertainty reduction for 10-bed ETC:

Timsteps after placement decision	1	2	3	4	5	6	7	8
% of uncertainty remaining	100	95	70	55	45	37.5	32.5	30

Function for the reduction in uncertainty of the transmission rate:

$$\%\ of\ uncertainty\ remaining = e^{-n/15} \times 100$$

where n is the cumulative number of patients. If $n > 100$, uncertainty is reduced to 0%.

C

Simulation Model Objectives

In order to calculate the Effectiveness objective, the epidemiological model is run a second time to generate the number of deaths that would have existed with no response. The number of deaths observed in the simulation of the actual response is then simply divided by this number, producing the percentage of people saved by the response. The efficiency ratio is also calculated with this percentage, by dividing it by the total costs accumulated over the response.

$$Effectiveness = \frac{deaths_{observed}}{deaths_{no\ response}}$$

$$Cost\ per\ Death\ Prevented = \frac{Effectiveness}{Total\ Cost}$$

The speed objective is implemented in the exact way it was formulated in Chapter 3 and requires not further explanation.

The two equity objectives are calculated as follows:
For Equity in Met Demand, the cummulative number of patients n_p is divided by the total number of cases n_c for each region i.

$$m_i = \frac{n_p}{n_c}$$

The average value of this ratio is then calculated, and the objective score is the sum over all regions R of the square differences between this average and the regional value for all regions. The lower the value of this sum, the better.

$$Difference\ in\ Met\ Demand = \sum_{i=0}^{R} (m_{avg} - m_i)^2$$

For Equity in Arrival Times, the onset of demand t_{onset} for each region i is taken as the first timestep at which more than 0 infections exist in that region. The difference between onset of demand and the time at which an ETC first becomes operational (t_{met}) is calculated for each region. If no ETC is operational before the end of the simulation the difference between onset of demand and the last timestep is taken. As for Equity in Met Demand, the average value over all regions is then calculated and the square differences are summed to produce the objective outcome value.

$$a_i = t_{met} - t_{onset}$$

$$Difference\ in\ Arrival\ Time = \sum_{i=0}^{R} = (a_{avg} - a_i)^2$$

D
Simulation Model Parametrization

The values of the model constants (taken from Büyüktahtakın et al. (2018) unless otherwise indicated) are:
$f_{without\ treatment}$: 0.31
$f_{with\ treatment}$: 0.24
$t_{death\ without\ treatment}$: $2.5(weeks)$
$t_{recovery\ without\ treatment}$: 2.85
$t_{death\ with\ treatment}$: 2.5
$t_{recovery\ with\ treatment}$: 2.37
$r_{safe\ burial}$: 0.64(Estimated from Rivers et al. (2014))
β_d : 0.73

Where f are the fatalitiy rates, t is the time in weeks until recovery or death, $r_{safe\ burial}$ is the safe burial rate and β_d is the transmission rate at a traditional funeral.

All regions except region 4, 14, and 15 start with 0 initial cases. For the number of susceptible people in each region, population data of Sierra Leone from 2015 is used (Statistics Sierra Leone, n.d.). The 4-by-4 grid is initialized with this data, with each grid cell corresponding to a district. During the 2014 epidemic, Sierra Leone had 14 districts, so two larger districts are split. This is purely done due to the practical requirement of having a square grid. The population data are shown in Table D.1.

No.	Region Name	Population
0	Kambia	345.474
1	Bombali #1	303.272
2	Koinadugu #1	204.686
3	Koinadugu #2	204.686
4	Porto Loko	615.376
5	Bombali #2	303.272
6	Tonkolili	531.435
7	Kono	506.100
8	Western Area (Urban)	1.055.964
9	Western Area (Rural)	444.270
10	Moyamba	318.588
11	Bo	575.478
12	Bonthe	200.781
13	Pujehun	346.461
14	Kenema	609.893
15	Kailahun	526.379

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15

Table D.1: Table showing the initial number of susceptible individuals in each region.

E

Validation of the Ensemble Size

All simulations were run with 2500 scenarios and 50 replications per scenario. In order to confirm that this ensemble of scenarios is large enough to draw valid results from in terms of model behaviour, convergence of a sensitivity analysis test was tested as suggested by Pianosi et al. (2016). As Difference in Arrival Time was found to be sensitive to several factors, and given that it showed a wide distribution of results, this objective was chosen for the test. The motivation for this choice is that if convergence is shown for this objective, it will also be present in objectives which are strongly influenced by only one or a few factors, which should lead to quicker convergence.

Feature scoring was carried out with the outcomes of the all-exploitation policy for arrival times below 850, the results of which are shown in Figure E.1. Between 500 and 2000 scenarios the results seem stable, though variation in the feature scores occurs again after 2000 scenarios. Still, the ordering of the of the factor influence is stable, which is considered essential Though full convergence would be ideal, it is concluded that an ensemble size of 2500 scenarios leads to acceptable results, especially considering that a more volatile objective was chosen as a test.

The all-exploitation policy is deterministic, but the policies defined by the RBFs have a stochastic element as they provide the probability with which an explorative action should be taken. This could result in more variation in outcomes and therefore ensemble size should be checked for these policies as well.The 141-policy was chosen for testing as the shape of its policy function would result in high levels of stochastic behaviour early in runtime, at which point it is believed to have the most impact on objective scores. Using the results of the 141-policy, convergence and stability of the ordering of factors were confirmed using the same conditions as for the all-exploitation policy.

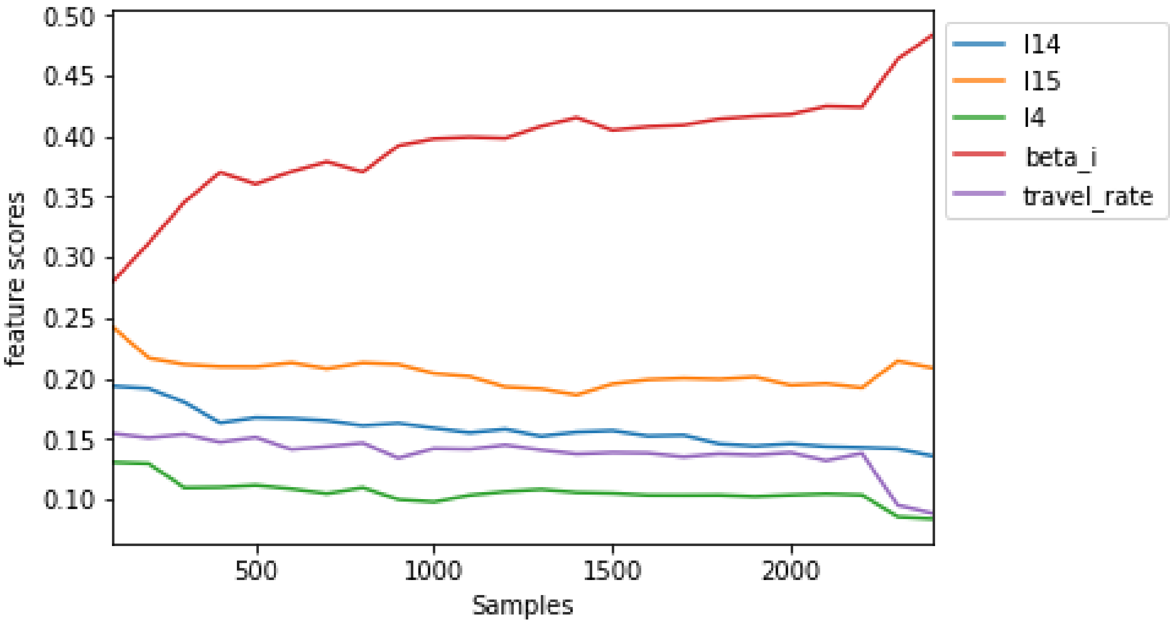

Figure E.1: Plot showing the convergence of feature scoring values for each of the objectives for the all-exploitation policy

F
MOEA Outcomes

As described in Chapter 7, two different Many-objective Evolutionary Algorithms (MOEAs) By comparing the Pareto fronts of the outcome objectives generated by both algorithms, and by comparing the policy shapes of the policies, we can get an indication of the quality of optimizations, considering neither were run until convergence.

The objective scores of the policies found by the Borg MOEA as well as the ϵ-NSGA2 algorithm are plotted in Figure F.1 [1]. At first glance, the outcomes from the Borg MOEA seem of higher quality - they lie closer to the bottom-left corner. However, it should be kept in mind that the ϵ-NSGA2 algorithm was run using 10 replications per scenario. For the Borg solutions, a policy may have received a higher score by chance due to the stochasticity of the model, but this is less likely to be the case for the ϵ-NSGA2 solutions as a result of the replications. Additionally, the ϵ-NSGA MOEA was run with larger ϵ values, which may result in better solutions not registering as improvements. Therefore, this difference in position can have two causes: a lower objective score due to the averaging over replications, as well as the actual quality of the solution.

For judging the quality of the solutions, it is also relevant to look at the shape of the Pareto fronts rather than their exact value when comparing the two results. Doing so, we can see that for most objectives, the Pareto Fronts have similar shapes, though they are more complete for the Borg solutions (i.e. more extensive and/or with fewer gaps). Only for the objective *Time until Containment* the solutions score notably different, which shows the effect of the replications. The outcome distribution of the Borg solutions shows two distinctive peaks, whereas the distribution of the ϵ-NSGA2 shows one peak, lying between those two peaks. This difference in distributions is caused by the averaging of outcome scores over the replications for the ϵ-NSGA2 solutions.

In comparing the results of the two MOEAs, it is also helpful to compare the policy functions associated with it solutions. For each objective, the policy functions associated with the best performance for that particular objective are shown side by side in Figure F.2. For the *Effectiveness*, *Time until Containment*, and *Cost per Death Prevented* the policy functions found by Borg and ϵ-NSGA2 are very similar, and the same variation in policies for a single objective is also seen. However, for the two objectives related to equity there are differences in the policy shapes provided by the two MOEAs. For *Difference in Met Demand* the same general bowl-like shape is visible, but in the policies functions of the ϵ-NSGA2 solutions the shape is more narrow. For *Difference in Arrival Time*, the ϵ-NSGA2 solutions all have a very distinctive dip towards exploitative actions near an uncertainty level just under 70%. In the Borg policy functions, this shape is only visible in one solution.

Several possible explanations exists for these differences: As a result of the lack of replications, some of the solutions found by the Borg MOEA may be “lucky” cases that resulted from fortunate

[1]These plots were generated in a Jupyter notebook called `MOEA Results` which is available at `https://github.com/edenbrok/thesis`

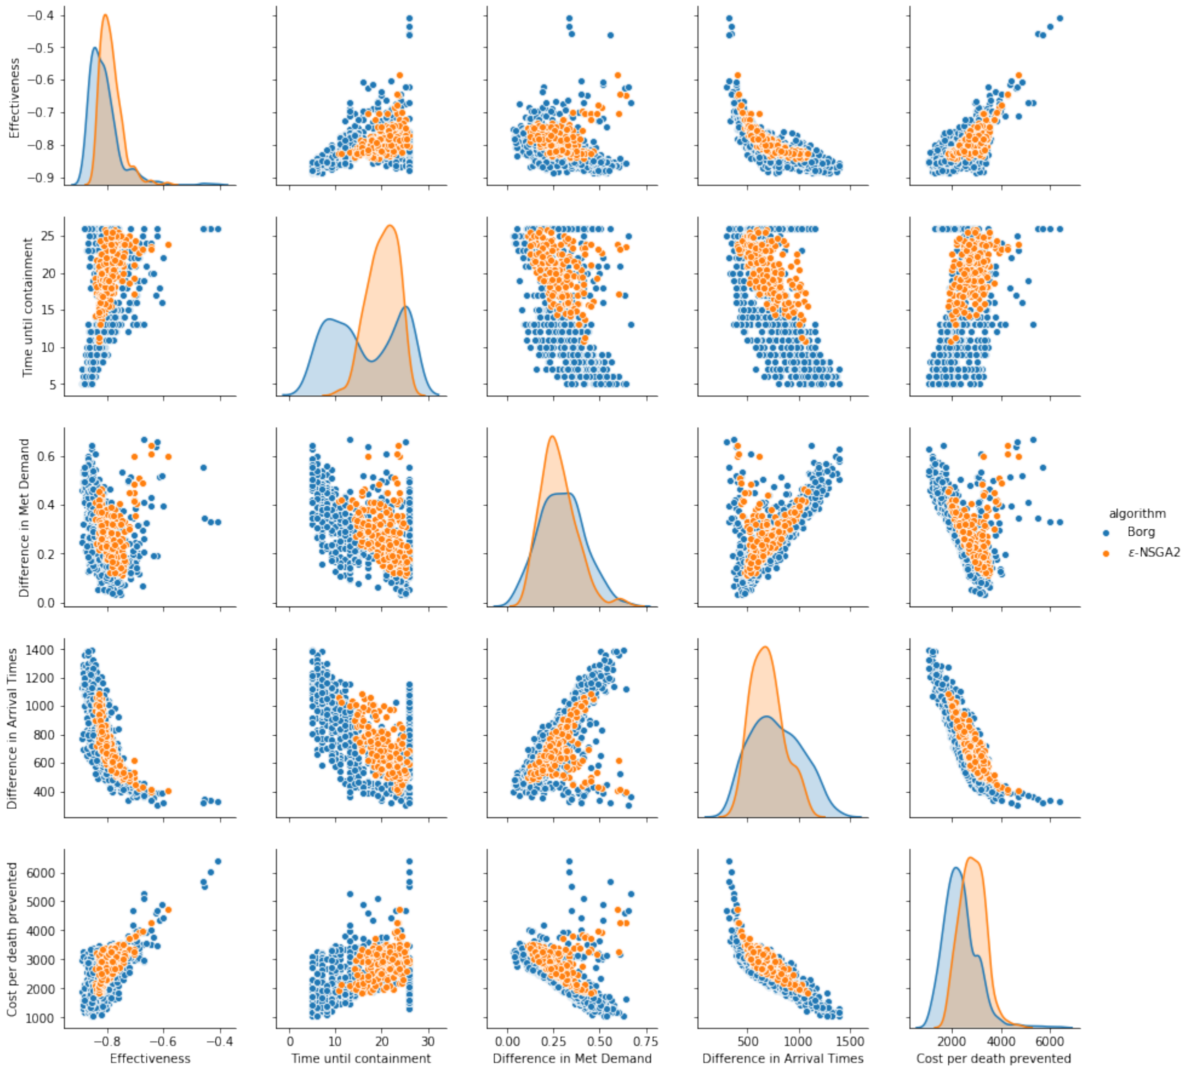

Figure F.1: Scatterplot showing the objective scores of the policies found by the Borg MOEA (blue) and the ϵ-NSGA2 algorithm (orange).

stochastic behaviour. However, these different policy shapes could also represent a family of solutions that had not been discovered by the ϵ-NSGA2 MOEA. Or, in the case of the policies for Difference in Met Demand, the policies found by ϵ-NSGA2 may not have been fully optimized yet, and would have resembled the Borg policies more closely given a higher number of nfe. Regardless, without further experimentation it is impossible to determine what the cause is and which policies are of higher quality.

Due to limited access to the compuational power necessary to run the ϵ-NSGA2 MOEA with replications, combined with time constraints, policy selections had to be made before the ϵ-NSGA2 outcomes were available and as a result were only based on the Borg outcomes. One policy representing the best performance on an objective was chosen - if scores were very similar, a policy with good scores on the other objectives was picked. The selected policies are shown in Figure 8.10.

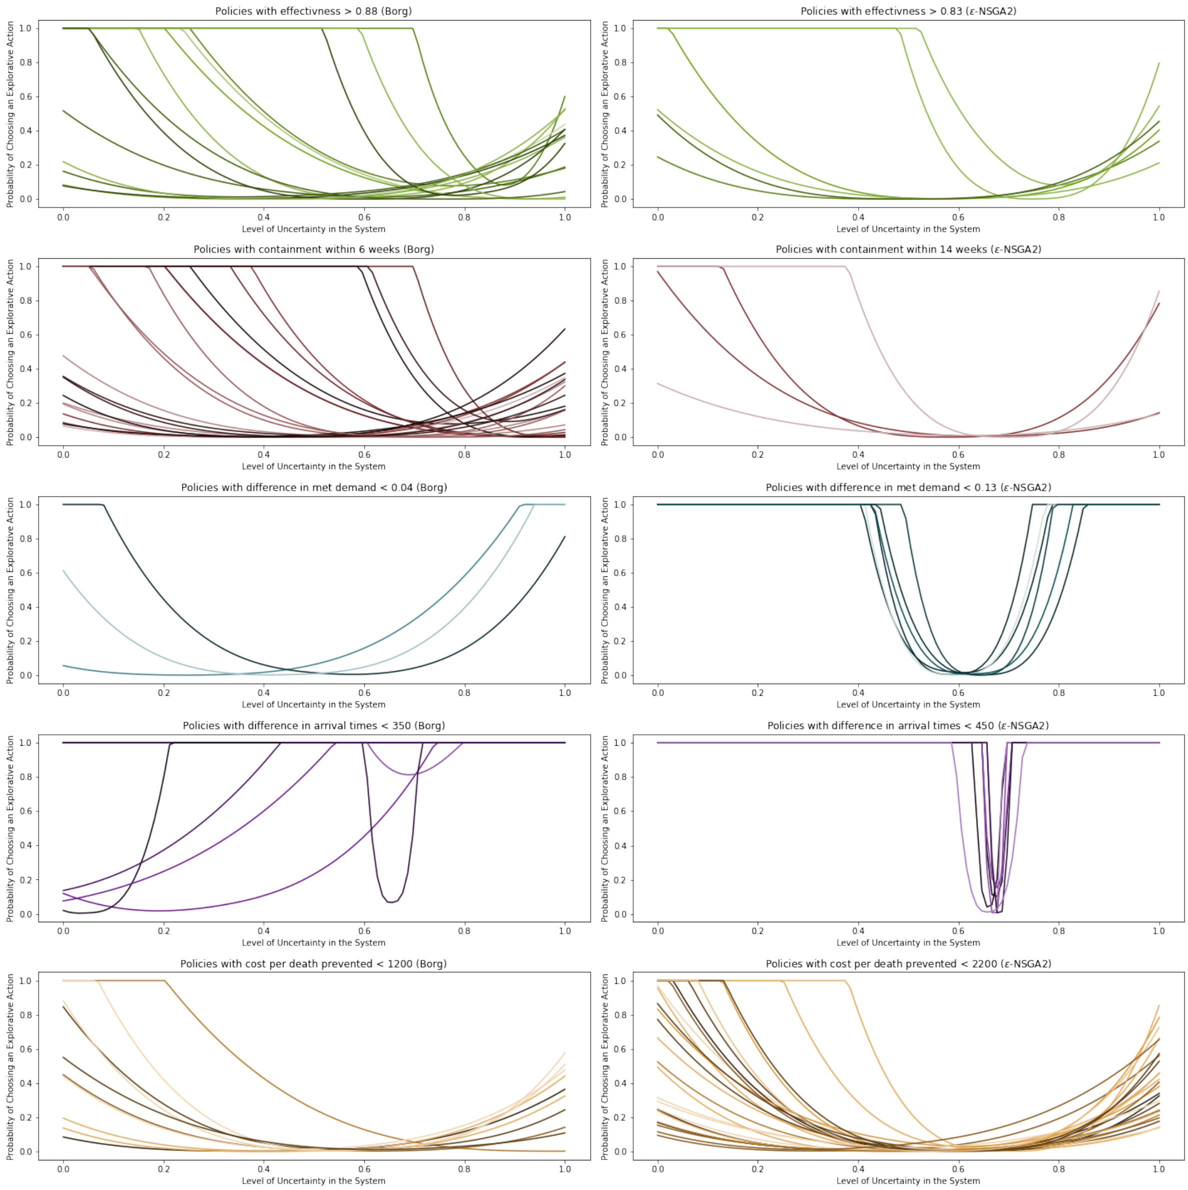

Figure F.2: Best performing policy functions for each objective as found by the Borg MOEA (left) and the ϵ-NSGA2 MOEA (right). Note that the selection criteria for the ϵ-NSGA2 outcomes are more lenient than those of the Borg outcomes, due to the differences in outcome scores discussed in the section above.

G
Scatterplots of All Policies

The following pages contain all the scatterplots of the policies discussed in this thesis.

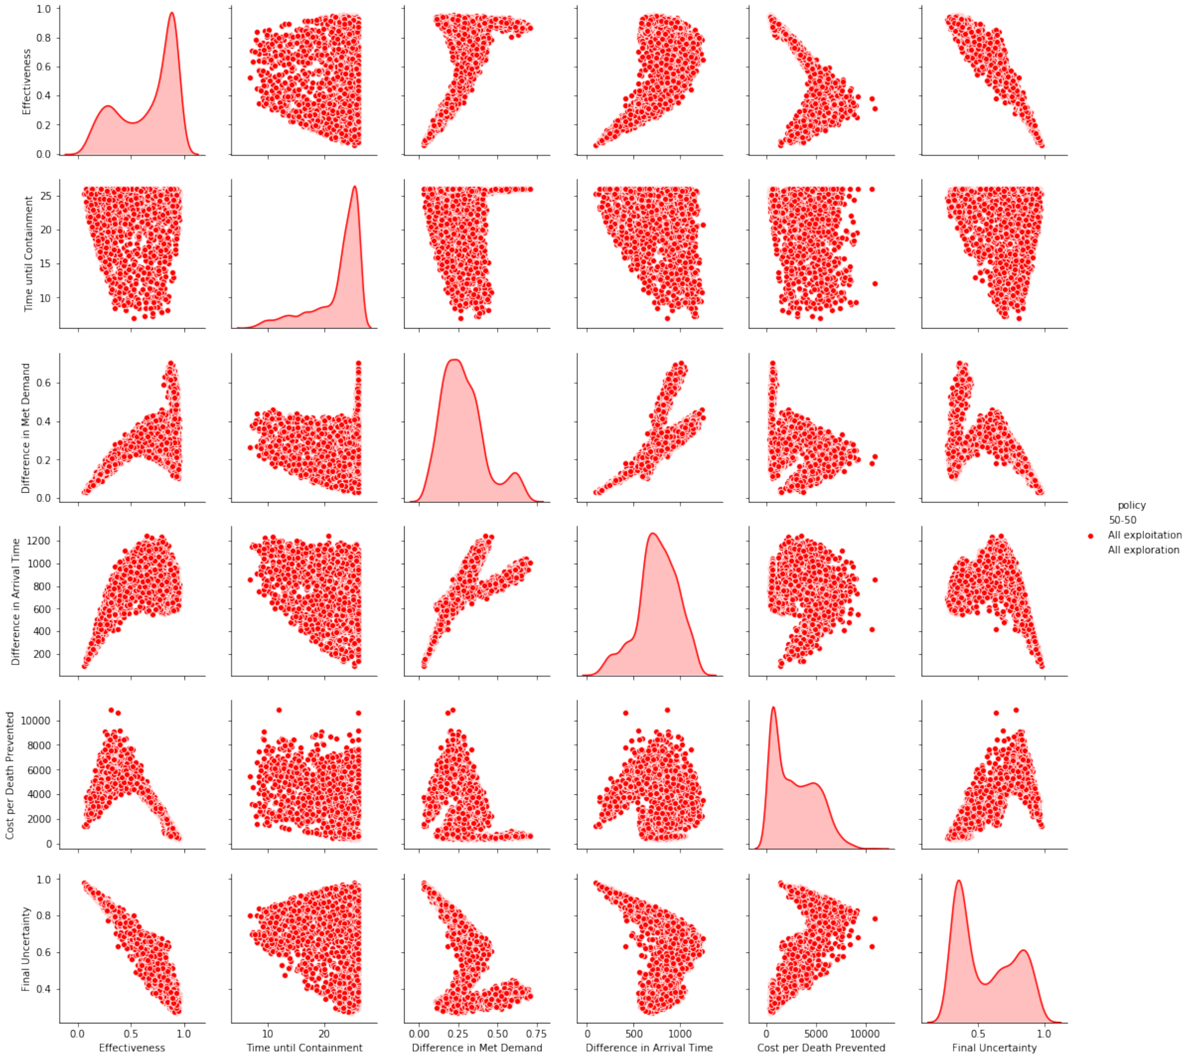

Figure G.1: Scatterplot showing the results of the all-exploitation policy for each objective plotted against all other objectives.

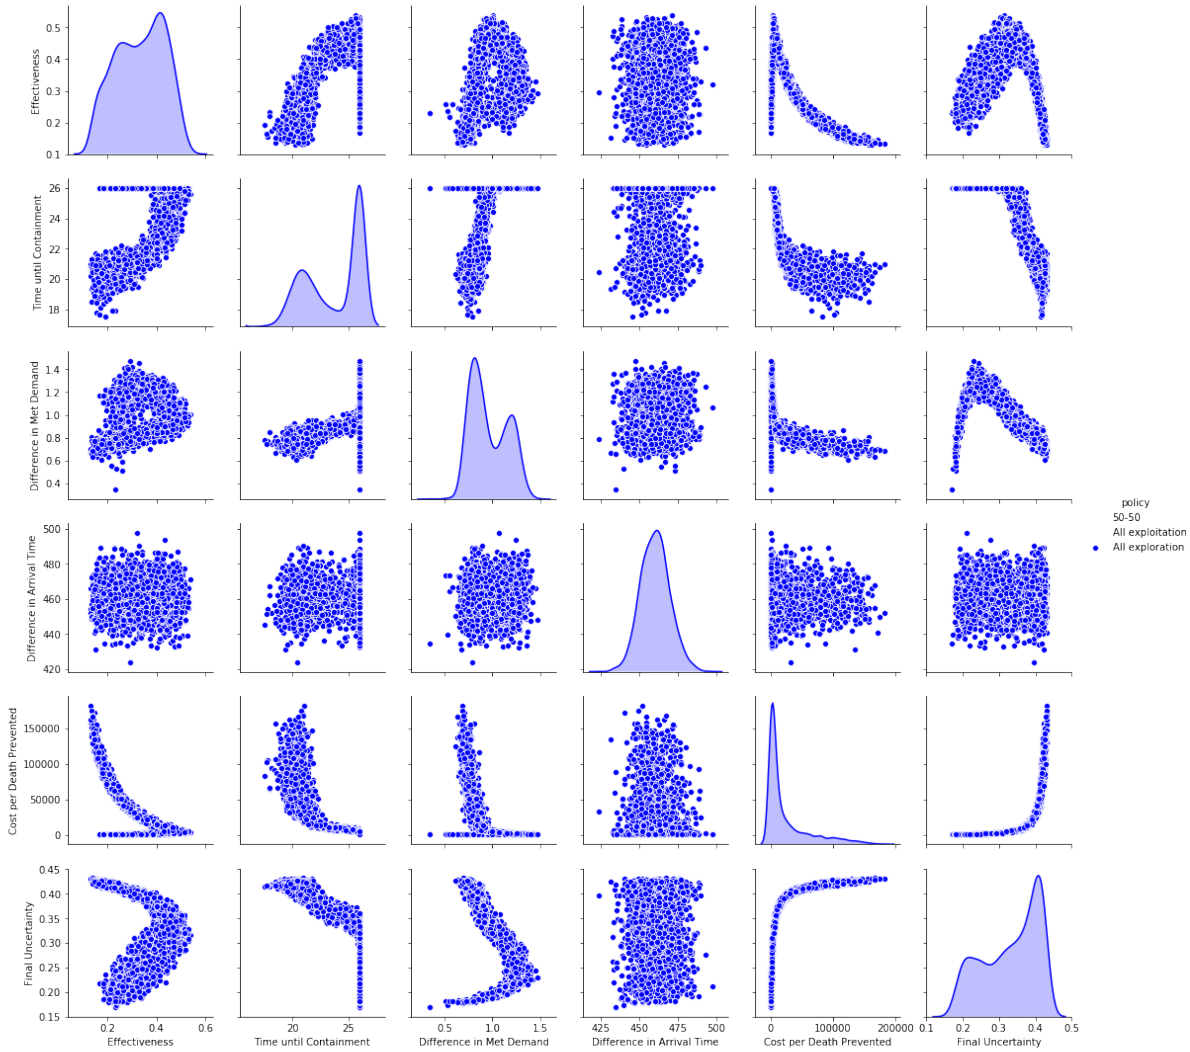

Figure G.2: Scatterplot showing the results of the all-exploration policy for each objective plotted against all other objectives.

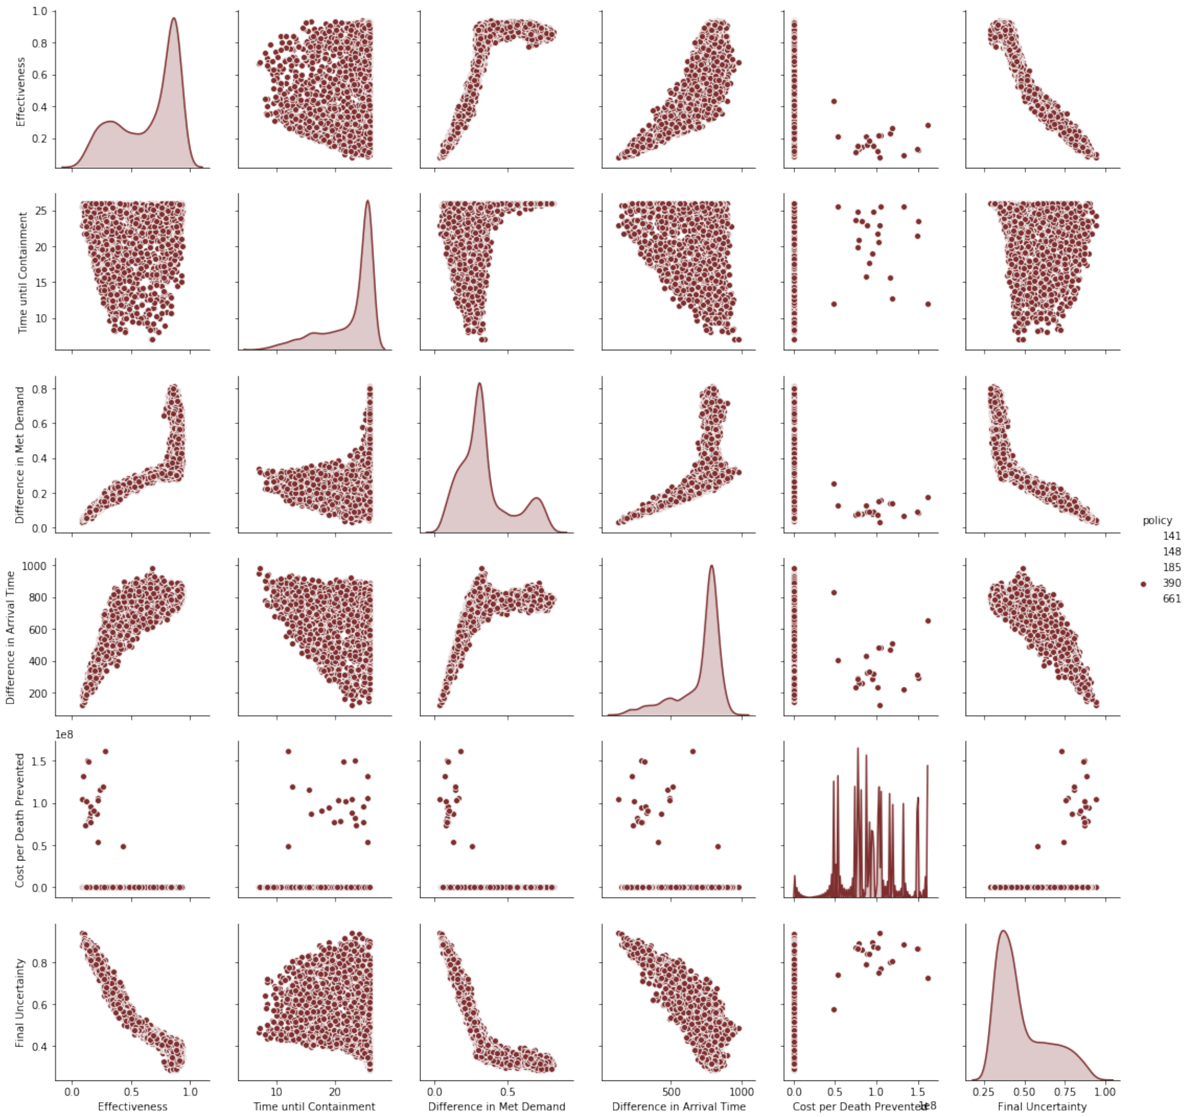

Figure G.3: Scatterplot showing the results of the exploitation-390 policy for each objective plotted against all other objectives.

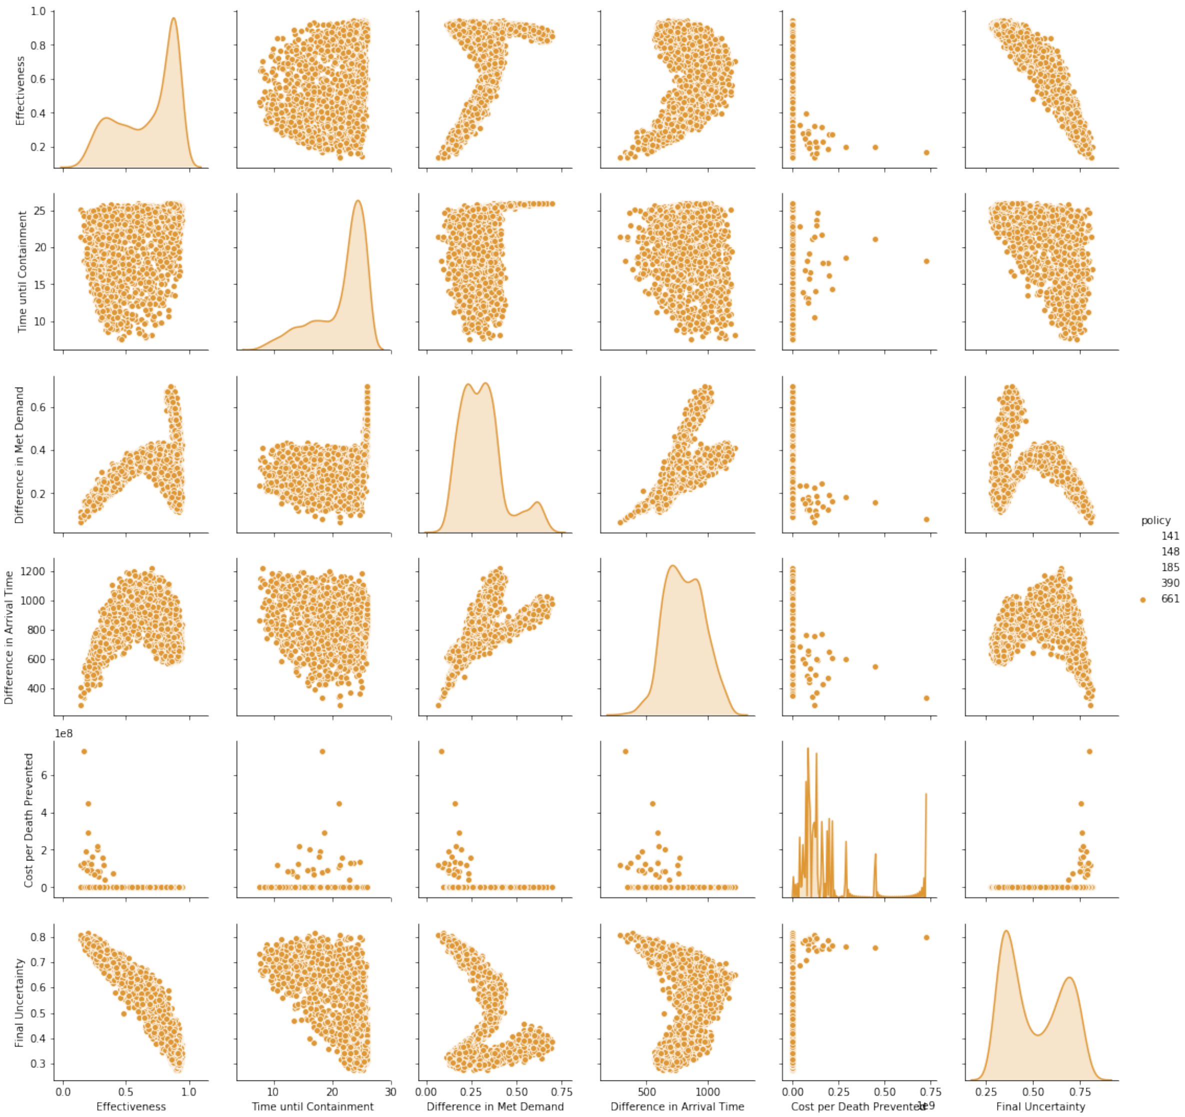

Figure G.4: Scatterplot showing the results of the exploitation-661 policy for each objective plotted against all other objectives.

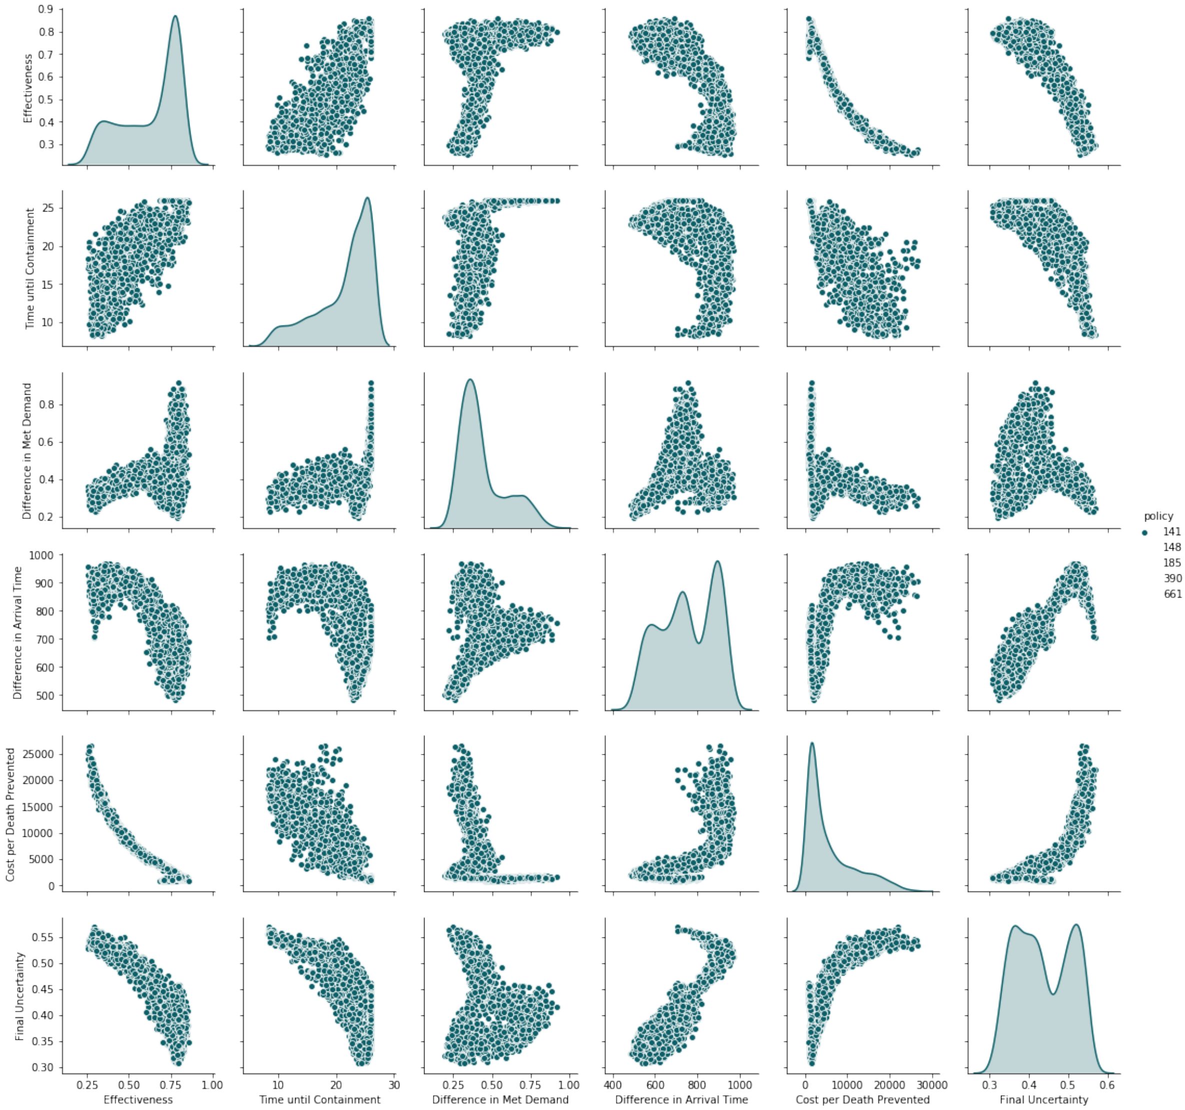

Figure G.5: Scatterplot showing the results of the exploration-141 policy for each objective plotted against all other objectives.

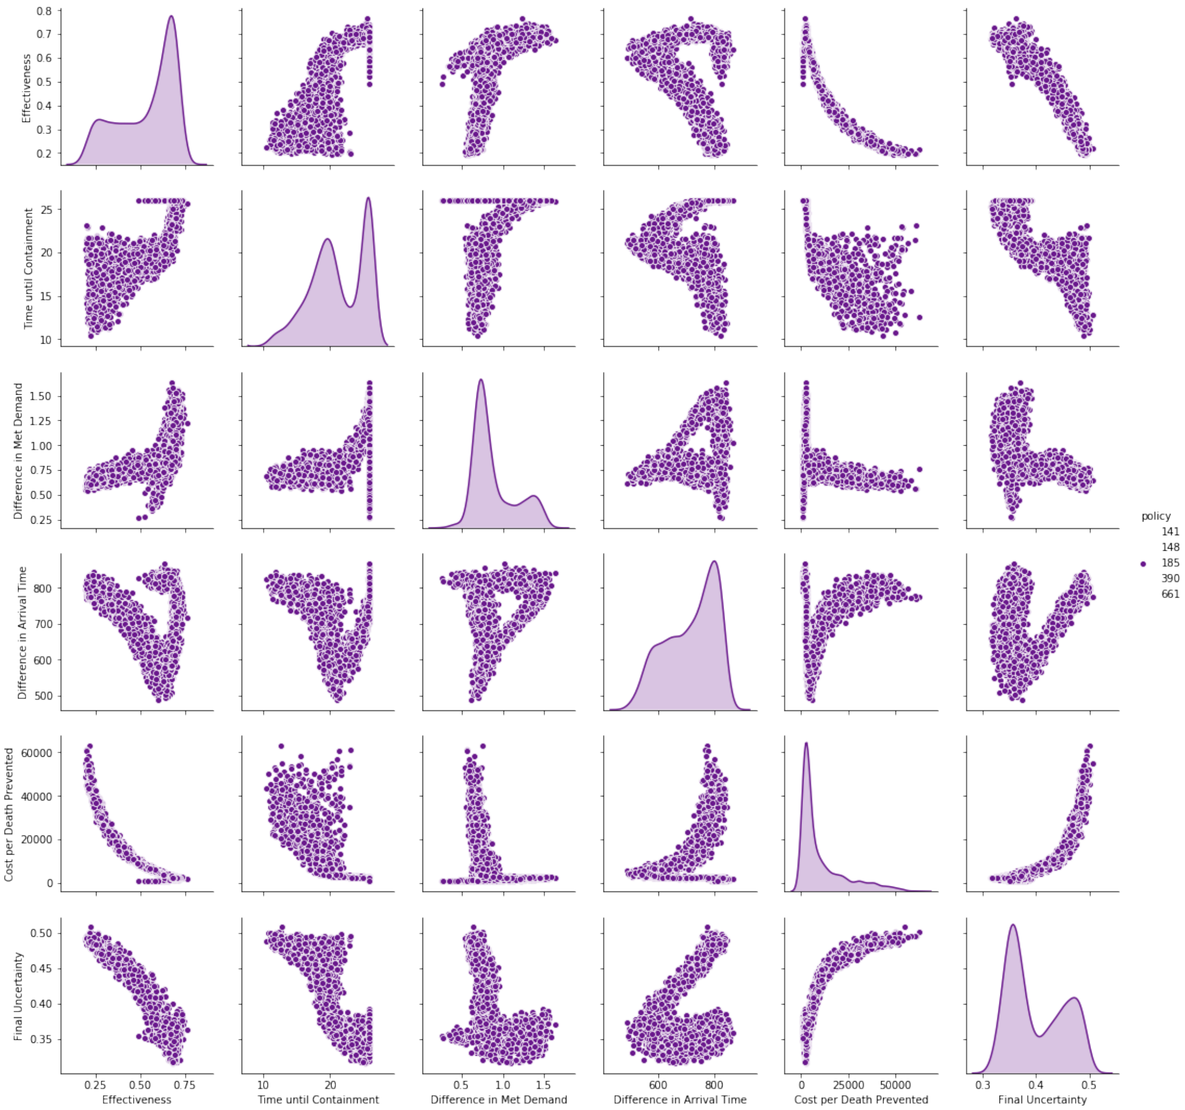

Figure G.6: Scatterplot showing the results of the exploration-185 policy for each objective plotted against all other objectives.

Figure G.7: Scatterplot showing the results of the random policy for each objective plotted against all other objectives.

H

Runtime Behaviour

H.1. Exploitation-390 Policy

The `exploitation-390 policy` starts out fully exploitative, but moves to a fully explorative policy at lower uncertainty levels. During the run pictured in Figure H.1, it receives spontaneous news from region 15 at timestep 2. By timestep 4, an ETC in that region is operational and the situation in the neighbouring regions is revealed. As uncertainty remains high, the response remains focussed on the hotspot regions (and their neighbours) for quite some time. However, after timestep 11 the policy starts exploring. By timestep 14 the decision-maker has a accurate picture of the situation, but as the policy is now fully explorative, this information is not used efficiently. However, since the main hotspot regions have been addressed early on in the response, it is still possible to contain the epidemic with the random placement of ETCs for explorative policies.

H.2. Exploration-185 Policy

Like exploration-141, the `exploration-185 policy` starts out fully explorative, before moving to a mixed and then fully exploitative state. However, the exploration-185 policy remains fully explorative longer and only starts to move to mixed decisions at a lower level of observed uncertainty. Figure H.2 shows snapshots of the runtime behaviour of the exploration-185 policy. Due to its explorative nature, it obtains a good situational overview early on (at timestep 3 it has already discovered all the hotspot regions (as spontaneous news from region 15 also came in)). However, it takes until timestep 6, when all regions have been discovered, for the level of uncertainty to be low enough for the policy to make an exploitative decision. For the remainder of the simulation run, the policy continues to take decisions of both types. The exploitative decisions can be targeted well due to the low levels of uncertainty, but because resources are also spent in a much more random manner by the explorative decisions, the policy fails to fully eradicate the epidemic.

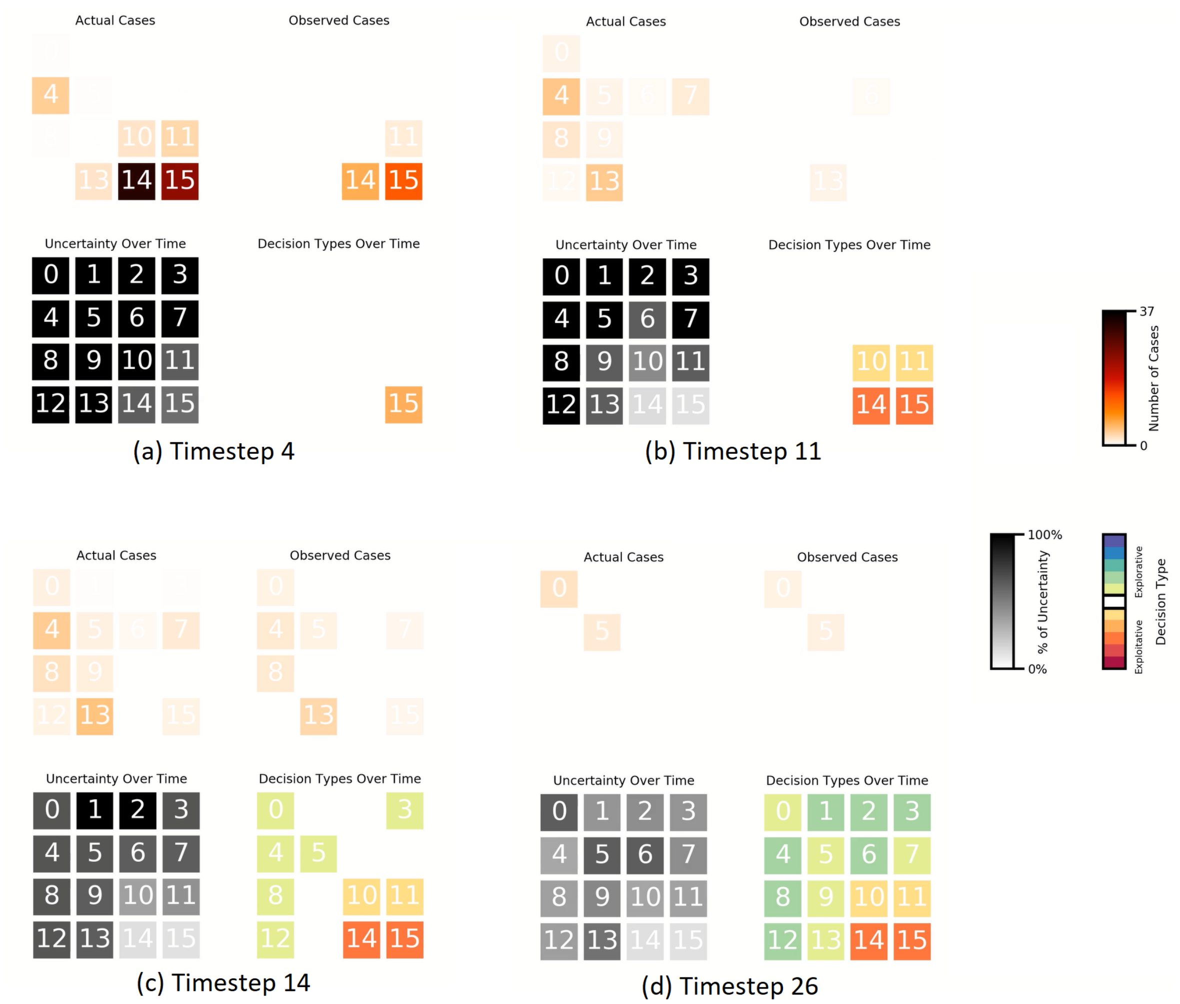

Figure H.1: State of the simulation model under the exploitation-390 policy and the base scenario at timesteps 4 (a), 11 (b), 14 (c) and 26 (d).

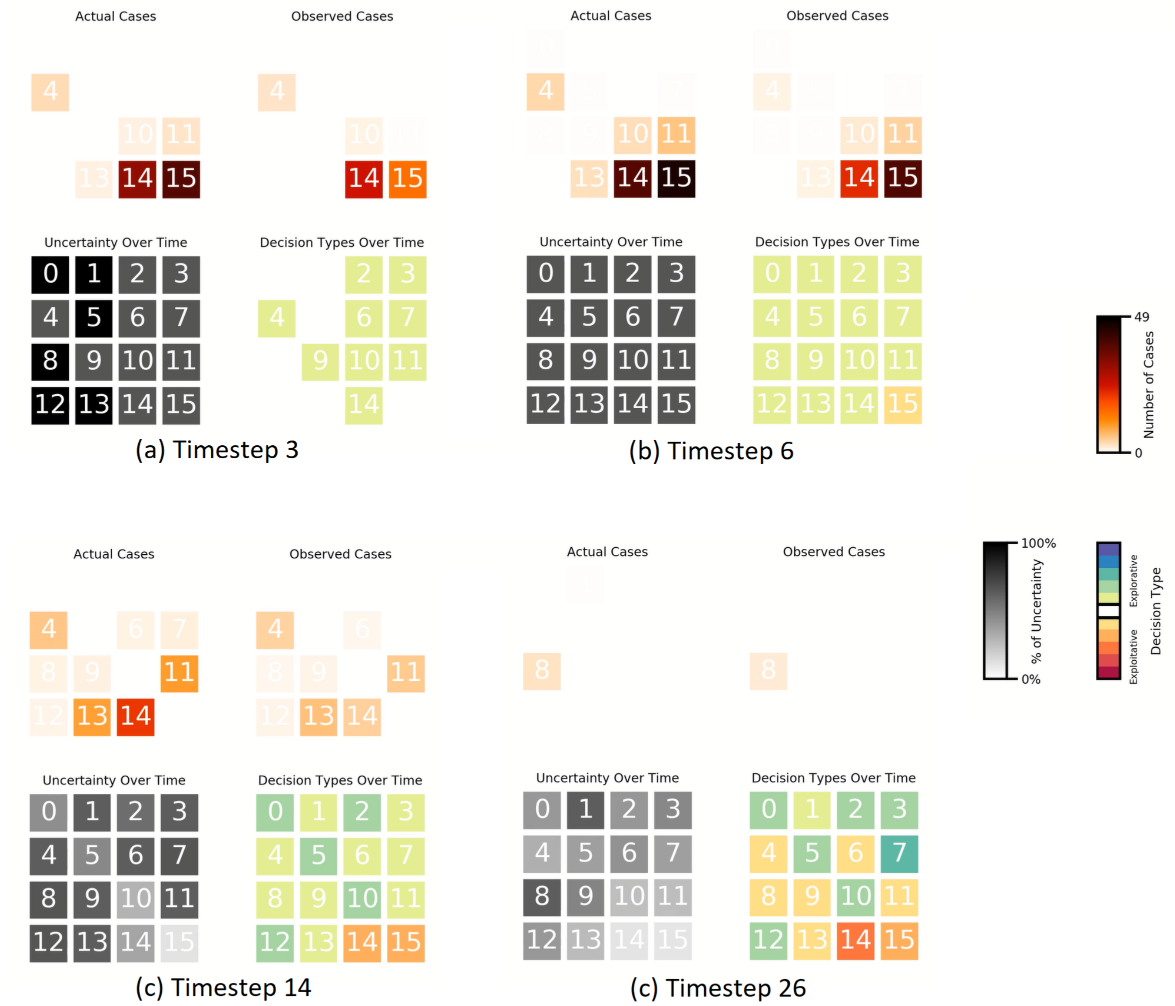

Figure H.2: State of the simulation model under the exploration-185 policy and the base scenario at timesteps 3 (a), 6 (b), 14 (c) and 26 (d).

I Code - Overview

All code is available on github at `https://github.com/edenbrok/thesis`.

`model_with_policy.py` This is the main file with the model function `borg_ebola`, which runs the simulation model based on the policy function parametrisation provided.

`model.py` Also allows for the simulation model to be run but with a constant exploration ratio using the function `ebola_model`.

`compartmental_model.py` Includes the `calc_population` function which is called every timestep by the simulation model to run the epidemiological model for 1 timestep. Also contains several helper functions for the compartmental model itself.

`decision_making.py` This file contains all the major functions used in decision making. `policy_exploration_ratio` calculates the exploration ratio using the policy function as described in 5.3.3. `explorative_decision` and `exploitative_decision` handle making their respective decision types. This file also includes functions for checking the number of resources in use, as well as functions that check if resources can be removed.

`uncertainty_reduction.py` Contains the functions `unc_infected` and `unc_transmission` that handle uncertainty reduction for each region as described in Appendix B. Also includes the function `total_uncertainty` that calculates the total level of uncertainty as experienced by the decision-maker.

`objective_functions.py` Provides the functions that calculate the model objectives at the end of each simulation run as described in Appendix C.

`objects.py` Contains all the objects used in the simulation model. `Uncertain_Constant` deals with the uncertainty experienced by the decision-maker for a constant factor (in this case the transmission rate). It keeps track of the ground truth as well as the range that is available for the decision-maker. `Uncertain_Variable` does the same but can handle variables whose ground truth changes during the simulation run, as is the case for the number of infected individuals in each region.
The `Region` class keeps track of all the data of each individual region; the values of each of the compartments over time, the uncertainty level, whether the region is hidden, the ETC capacity etc. It has functions to update itself after the compartmental model is run or after resources have been placed in a region. It also contains the “spontaneous news” function.

The `ETC` and `Surveillance_Team` classes are used to keep track of the timesteps in which the resources are used, operational and/or closed so that the uncertainty reduction and region updating functions can be used easily.

`utility.py` Contains two functions: The first is `get_neighbours` which returns a list of neighbouring regions when it is fed a square grid and a region index for which the neighbours need to be found. This function is used by the compartmental model to determine the influx of infected people from neighbouring regions. It is also used by the other function in this file, `random_travelling`, to ensure that the superspreading travelling events go from one region to a non-neighbouring region.

J Interviews

All interviews were conducted over Skype, recorded, and transcribed. Transcripts were sent to interviewees for approval. Anonymised transcripts for each interview can be found at `https://github.com/edenbrok/thesis/tree/master/Interviews`.