Three essays on decision making in health care operations

Eike Nohdurft

Dissertation in partial fulfillment of the requirements for the degree of
Doctor of Economic Sciences (Dr. rer. pol.)

WHU – Otto Beisheim School of Management
Vallendar

February 26th, 2017

Primary Supervisor:
Prof. Dr. Stefan Spinler

Second Supervisor:
Prof. Dr. Arnd Huchzermeier

Abstract

In this thesis, we[1] use operations research methods to provide insights into three areas associated with health care operations management. In Chapter 2, we use a discrete-event supply chain simulation to asses if coordination among partners is beneficial in a supply chain with the characteristics of the German pharmaceutical market. We find that the greatest cost savings and service levels could be achieved through a highly integrated collaboration although most of its impact could already be achieved through sharing point-of-sales demand information. Results suggest that coordination is most beneficial in situations where product shelf life is short and demand variation is high.

In Chapter 3 we consider quality-of-life maximizing sequences of prophylactic surgeries for female carriers of a BRCA1/2 genetic mutation, who face a significantly elevated breast and ovarian cancer risk. Using a Markov Decision Process model, we determine the optimal surgery sequence that maximizes the carrier's expected lifetime quality-adjusted life years (QALYs). Baseline results demonstrate that a QALY-maximizing sequence recommends a bilateral mastectomy between ages 30 and 60 and bilateral salpingo-oophorectomy after age 40 for BRCA1 carriers. Surgeries are recommended later for BRCA2 carriers, as their cancer risk is lower. The model's structural properties show that when one surgery has already been completed, there exists an optimal control limit after which performing the other surgery is always QALY-maximizing.

In Chapter 4, we develop a two-stage model for optimizing when and where to assign Ebola treatment unit (ETU) beds—across geographic regions—during an infectious disease outbreak's early phase. The first stage includes a dynamic transmission model that forecasts occurrence of new cases at the regional level, thus capturing connectivity among regions; in this stage we introduce a coefficient for behavioral adaptation to changing epidemic conditions. The second stage includes two approaches to efficiently allocate intervention resources across affected regions. Such an allocation could have prevented up to 3,434 infections over an 18-week period during the 2014 Ebola outbreak in West Africa, a 58% improvement compared with the actual allocation.

[1] In Chapter 2, 3, and 4, the term 'we' refers to the authors of Nohdurft & Spinler (2016), Nohdurft et al. (2016a), and Nohdurft et al. (2016b), respectively.

Acknowledgements

During the past months I've been looking forward to the moment when I would write this brief, albeit important, section of my thesis. Many people have contributed to the fact that I've gotten this far, everyone in his or her unique way.

My main advisor, Prof. Dr. Stefan Spinler, has put a tremendous amount of trust and support in me from day one. He not only encouraged a very preference-based selection of research areas, but also had invaluable impact on all projects through his precise methodological advice. Stefan never hesitated when something was needed, were it resources or flexibility. The second outstanding person in this endeavor has been Prof. Elisa F. Long, Assistant Professor at UCLA Anderson School of Management. Elisa accepted me as a visiting scholar in Los Angeles while I had almost no research credentials and invited me to work on BRCA, a project based on a very personal experience of hers. She is a very inspiring co-author and discussion partner who, together with her husband Keith, gave me insights into the American culture which I though I already knew. I also would like to thank Prof. Dr. Arnd Huchzermeier, my second advisor, for taking the time to review this manuscript.

None of this would have been possible without the support from my parents, my grandparent, my sister and her family. They did not doubt once if I'm doing the right thing. While every Ph.D. student knows that for every high there's at least one low already waiting, since a particular exam supervision in 2014 I'm having the fortune of being able to share all highs and lows together with Inga. I also would like to thank my long-time friends in Hamburg for being there when I needed their help and Michael, the most rigorous discussion partner when it came to LaTeX syntax.

When writing a Ph.D. in Operations Research, one inevitable spends a lot of time in an office in front of a computer. I therefore would like to thank my colleagues at the chair for the countless coffee breaks, discussions, favors, and jokes.

Contents

List of figures

List of figures

List of tables

List of abbreviations

ALL	allocation	26
ADP	Approximate Dynamic Programming	71
ANOVA	analysis of variance	27
BM	bilateral mastectomy	38
BSO	bilateral salpingo-oophorectomy	38
CAP	capactiy	26
CSL	customer service level	11
CONV	conventional supply chain information sharing	19
CPFR	collaborative planning, forecasting and replenishment	12
DP	Dynamic Programming	20
FCFS	first-come-first-serve	17
FR	fill rate	21
HR	hazard ratio	116
IFR	increasing failure rate	50
LP	linear programming	39
LuA	Luminal A	44
LuB	Luminal B	44
MANOVA	Multivariate Analysis of Variance	26
MDP	Markov Decision Process	39
MRI	Magnetic Resonance Imaging	60
POS	point of sale information sharing	19
PRAC	practical allocation	25
PROD	product	26
PROP	proportional allocation	25
QALY	quality-adjusted life year	39
QOL	quality-of-life	38
OR	odds ratio	116
SC	supply chain	10

SCC	supply chain costs	21
SCS	supply chain scheme	26
SL	shelf life	26
VMI	vendor managed inventory	11
WHO	World Health Organization	2
TN	Triple-Negative	44

Chapter 1

Introduction

1.1 The need for change in the health care economy

The global health care economy is facing three substantial developments. First, health care systems in developed countries are continuously increasing their level of care following a never-ending evolution of diagnosis, treatment, and medication options. While this imposes a heavy cost increase on health care providers and payors, i.e. through cancer treatments costing more than 100,000 US$ per year (Siddiqui & Rajkumar, 2012), patients also secondly face more complex decisions when being asked to choose from a wider set of possible diagnostic or surgical procedures. Third, especially in emerging and developing countries, natural disasters and infectious disease outbreaks pose a major challenge for the local health care systems, as they are seldom prepared for large-scale health crises.

As a reaction to increasing cost of care, payors regulate the prices of newly introduced pharmaceuticals, i.e. in Germany through a reference price scheme (Henschke et al., 2013). Combined with increasing cost for research and development of new high-revenue pharmaceuticals, pharmaceutical manufacturers are starting to consider increasing the efficiency of their operations after decades of ever-high profit margins (Ward, 2016).

In most health-affecting situations, patients can rely on recommendations given by their physician. As diagnostic procedures analyzing genetic information become cheaper, a decision whether to perform these tests and what actions to conclude from their re-

sults often also depends on the patient's evaluation of the situation (Dancey et al., 2012). While physicians can still provide recommendations and counseling, solutions transforming the complex stochastic relationships into easy-to-understand recommendations are needed (Reyna et al., 2015).

While in developed countries, health care systems are trying to keep cost low while care quality is increasing, global health crisis caused either by natural disasters or infectious diseases are hitting developing nations with increasing frequency and impact. Natural disasters, like wildfires or earthquakes do only have a local or regional impact (e.g. the 2015 earthquake in Nepal). Infectious disease epidemics on the contrary can cross borders and even oceans in few days or even hours through the international airline network, converting a local infectious disease epidemic rapidly into a global health concern (Mangili & Gendreau, 2012).

As local health authorities are often unable to contain an infectious disease outbreak once it peaks, foreign government and non government organizations provide monetary, infrastructure, or personnel support (United Nations, 2015; World Health Organization, 2003). With many stakeholders involved, coordination of efforts poses a major challenge for institutions as the World Health Organization (WHO) (Philips & Markham, 2014). Complexity is further increased through the nonlinear behavior of an epidemic, which is determined by various factors like the disease prevalence, its transmission dynamics, or the behavior of infected and susceptible individuals. Today's literature lacks effective methods for coordinating and allocating resources during an infectious disease outbreak, which reduce the complexity of decisions and stakeholder coordination throughout the epidemic.

1.2 The pharmaceutical supply chain

While pharmaceuticals already represent a significant share of the total health spend in developed countries (i.e. 12% in the US or 14% in Germany, OECD, 2016), prices tend to increase further due to the high development and certification efforts required for bringing a new pharmaceutical to market (Loftus, 2016). Payors try to lower these cost through a more cost-sensitive certification process (Thomas, 2016). To maintain their margins, pharmaceutical manufacturers usually provide the highest cost savings by themselves, while supply chain cost also present a viable cost reduction

lever (Schneller & Smeltzer, 2006).

The pharmaceutical supply chain differentiates itself from other supply chains through transporting highly sensitive goods, which are object to temperature and shelf life regulations, to patients, who rely on timely delivery as otherwise they might face a negative, potentially life threatening, health impact. To cope with these challenges, the industry sticks to high service level standards through maintaining high inventories and providing frequent and fast delivery to the retail outlet, mostly pharmacies (Scheel et al., 2014). To overcome these cost-intensive policies, the pharmaceutical supply chain could adapt coordination as supply chain management method which has already reduced supply chain cost in other industries (de Vries et al., 2011). As evidence for the benefits of supply chain collaboration is not without doubt, we[1] aim to investigate the benefit specifically for the pharmaceutical supply chain and make the following contributions to the literature:

(i) We build a comprehensive discrete-event simulation model reflecting the characteristics of the German pharmaceutical supply chain (SC) to assess the impact of SC coordination on the SC cost and fill rate. The degree of coordination between the supply chain entities can be chosen among three different collaboration methods. Other parameters can be varied to reflect different environments in which the SC operates. These parameter include production capacity, product shelf life, and different order allocation mechanism used in the case of constrained production capacity.

(ii) The demand data used throughout our experiments is based on real-world demand data collected in the German pharmaceutical market. The dataset enables us to conduct experiments with demand patterns from different pharmaceuticals. For the use in our empirical experiments, we bootstrap the data using a maximum entropy bootstrap, a method which has not been previously used in the context of discrete-event simulation.

[1] 'We' in Chapter 2 refers to the authors of Nohdurft & Spinler (2016)

1.3 BRCA1/2 carrier decision support

An increase in genetic testing capabilities of the health care industry and a decrease in the involved cost has made genetic testing more widely available among affected and unaffected patients (Katsanis & Katsanis, 2013). Although the test procedures become increasingly precise, determining a diagnose, the corresponding treatment or prophylactic measures from the test results is not always straightforward. Additionally, patients are often not able to interpret the test results and the corresponding health risks adequately (Portnoy et al., 2010).

One of the genetic mutations for which a test is becoming increasingly less expensive is BRCA1/2. It has a prevalence of 1 out of 400 women in the U.S. – and one in 50 women of Ashkenazi Jewish descent (Hall et al., 2009). Female carriers have an increased lifetime breast and ovarian cancer risk of up to 85 and 54%, respectively, compared to 12 and 1.4% in the general population (King, 2003; Howlader et al., 2012). While the mutation prevalence is low in the general population, it varies across racial/ethnic groups and is higher in families with a history of breast cancer (Frank et al., 2002; Malone et al., 2006). Options to reduce the cancer risk include medication, e.g. tamoxifen therapy, or a prophylactic removal of the breasts and ovaries. While the use of tamoxifen has shown to moderately reduce the risk of a contralateral breast cancer (Gronwald et al., 2014), a prophylactic bilateral mastectomy reduces the breast cancer risk by about 91% and a bilateral oophorectomy the breast as well as ovarian cancer risk by up to 59 and 79% (Eisen et al., 2005; Rebbeck et al., 2004, 2009). While each surgery is an invasive procedures accompanied by potential complications and long-term physiological and psychological side effects, the literature reports heterogeneous quantitative results about the impact of prophylactic surgeries on carrier's quality-of life (Harmsen et al., 2015).

Although guidelines when to test for a BRCA1/2 mutation exist Moyer (2014), the medical decision making literature still lacks comprehensive recommendations if or when to perform prophylactic surgeries. We[2] therefore make the following contributions with the goal to provide a model supporting decisions about prophylactic surgeries with understandable and actionable recommendations:

(i) We build a comprehensive Markov Decision Process (MDP) which models cancer

[2] 'We' in Chapter 3 refers to the authors of Nohdurft et al. (2016b).

risk and mortality of BRCA1/2 carriers. The model is calibrated with data from clinical studies for each mutation type. Unlike previously published studies, we differentiate between four different tumor stages and breast cancer subtypes to capture the impact on the cancer-specific mortality rate. We also consider an age-dependent impact of a oophorectomy to account for the fact that losing the ability to reproduce has a higher impact on quality-of-life especially for young women.

(ii) Our numerical results provide an indication about the optimal timing of prophylactic surgeries maximizing a carrier's discounted cumulated lifetime quality-adjusted life years. Although the size of the MDP's state space is too large to be solved through policy or value iteration, we transform the problem to a linear program which we are able to solve optimally. Results include the optimal surgery sequence for carriers of different ages and for a variety of different model parameters during an extensive sensitivity analysis. We also provide results when the choice of possible surgeries is constrained through either personal preferences or an earlier cancer diagnosis.

(iii) We provide structural insights into a simplified version of the MDP and show that with only one surgery left in the choice set there exists an optimal control threshold for the cancer risk after which it is always optimal to perform the remaining surgery.

1.4 Resource allocation during infectious disease outbreaks

Infectious diseases have been, and albeit advances in their prevention and treatment, still are a major concern of health care systems around the globe (Fauci & Morens, 2012). Global air travel and an increase in population density in urban areas foster their emergence and re-emergence (Morens et al., 2004). Especially the concern over a global epidemic, a so-called pandemic, has increased (Lightfoot et al., 2013). Several outbreaks of infectious diseases in the last years (e.g. SARS in China, Ebola 2014 in West Africa, and Zika 2015 in South America) had at least the theoretical potential of causing a pandemic.

For the containment of a fast-growing epidemic, three major challenges have to be considered: (i) the amount of intervention resources required is quickly exceeding the capabilities of the affected health care system. (ii) Especially for highly contagious diseases, the future trajectory and spread of the outbreak is unknown and difficult to forecast precisely. (iii) When intervention resources are provided by foreign government or non-government organizations, their efficient coordination is often hampered through cross-organizational boundaries such as differing standard, hesitation to exchange data, or language barriers.

While during the 2014 Ebola outbreak in West Africa, sufficient monetary resources had been pledged by different organizations (World Health Organization, 2016), the effectiveness of the response was lowered through an initial overestimation of case numbers (Butler, 2014b) and the lack of coordination of intervention resources (Gettleman, 2014). Both challenges could have been faced through a better utilization of the available data on where and when infections occurred. While an accurate epidemic model could forecast the spatial spread and infection numbers of the outbreak, its outcome could also be used to increase the effectiveness of the allocation of intervention resourced across the afflicted regions.

Although the 2014 Ebola epidemic in West Africa has fortunately abated, we[3] utilize the available data on Ebola cases to make the following contributions to the existing literature body on epidemic modeling and intervention resource allocation during infectious disease outbreaks:

(i) We develop an epidemic model considering multiple, interdependent populations, which reflect the heterogeneity in Ebola prevalence among the afflicted regions. To incorporate the change of social behavior during an epidemic, we propose a novel dampening coefficient which adapts the speed of transmission to the change in social behavior of individuals (e.g. minimizing social contact) who live in the affected communities. When fitting the model to the data from the 2014 Ebola outbreak in West Africa, we find improved forecasting capability both for spatial resolution and fit quality compared to earlier models.

(iii) We provide an analytical approach for calculating the basic reproduction number, R_0, for our multi-population epidemic model. We use the numeric results of R_0 in an allocation heuristic and compare it with the results of the ADP

[3] 'We' in Chapter 4 refers to the authors of Nohdurft et al. (2016a)

algorithm.

(ii) We use a dynamic programming algorithm to determine the allocation of intervention resources, beds in our experiments, to the afflicted regions in Guinea, Liberia, and Sierra Leone. While the size and multi-dimensionality of the state space make the computation of an optimal solution intractable, we develop an approximate dynamic programming algorithm based on policy iteration which is able to solve the problem in reasonable time. The approach complements the heuristic using R_0 and both are compared to one other allocation heuristics as well as the actual allocation. It is, to the best of our knowledge, the first application of approximate dynamic programming with an underlying epidemic model.

Chapter 2

Coordination benefits in pharmaceutical supply chains

with Stefan Spinler

The content presented in this chapter is based on Nohdurft & Spinler (2016). Partial results from this chapter contributed to the conference presentations Nohdurft & Spinler (2014a) and Nohdurft & Spinler (2014b).

The health care supply chain in developed markets is a substantial cost driver in an industry under permanent cost-pressure while proven supply chain practices from other industries, like retail, have not been broadly applied. Coordination methods, like collaborative planning, forecasting, and replenishment (CPFR), could lower cost while maintaining stable performance levels. We develop a discrete event simulation which evaluates the impact of two coordination methods on supply chain performance considering different external constraints. Results show that coordination is beneficial both for supply chain costs and customer fill rate, especially when product shelf life is short. The model further indicates that implementation of a coordination scheme needs to be accompanied by a profit sharing agreement, as the benefits are highest for the wholesaler echelon of the supply chain.

2.1 Introduction

Modern supply chains (SCs), transporting goods from manufacturers to a diverse network of retailers across country and company borders, have become complex networks where performance has a significant financial impact on the profit of the involved stakeholders. The health care sector, which is seeking to enhance efficiency under ever-rising cost pressure, not only needs to handle such a complex network but is also further constrained through regulatory hurdles, sensitive goods, and high service level requirements. The benefits of intensified information sharing and coordination could help to to reduce the cost of a pharmaceutical supply chain while simultaneously maintaining a high performance level.

While the effects of coordination on SC performance have been evaluated in different settings (Aviv, 2001; Byrne & Heavey, 2006; Cachon & Fisher, 2000; Kamalapur et al., 2013), we contribute to the literature by focusing on the narrow but important market of pharmaceuticals. Accounting for about 30% of health care provider's overall cost (Schneller & Smeltzer, 2006), supply chain management has the potential to be an improvement lever for an industry facing high cost pressure while having to comply with high service levels (Burns, 2002; Dacosta-Claro, 2002). The handled goods are mostly perishable, which makes most of the existing research difficult to apply. Our proposed model simulates different degrees of coordination in a discrete event simulation which is calibrated to the German pharmaceutical market. We use real demand data to assess potential performance gains through the application of three coordination schemes. Results show that coordination is beneficial under most circumstances, but the size of its benefits varies with the shelf life, demand variation and available capacity.

The remainder of the paper is organized as follows: Section 2.2 provides an overview about existing research on SC coordination and performance of health care SCs. Section 2.3 describes the simulation model setup. Section 2.4 provides details about the design of the experiment, while section 2.5 analyses its output. The paper concludes with a summary of the results and provides possible directions for future research in section 2.6.

2.2 Literature review

The following review provides an overview about four different forms of SC coordination and SC evolution and performance in the health care industry.

Supply chain coordination The field of SC coordination is well represented in the literature[1]. Results from empirical as well as analytical studies shows, that it can have significant positive impact on cost and inventory level. Kanda & Deshmukh (2008) propose four different coordination mechanisms: supply chain contracts, information sharing, information technology and joint decision making.

Supply chain contracts regulate the legal relationship between SC partners with the objective to increase overall supply chain performance while simultaneously managing the risks emanating from coordination for each partner involved (Kanda & Deshmukh, 2008). The possibility of performance improvement is shown by Tsay (1999) through modeling the benefits of a contract regulating product prices as well as minimum and maximum purchasing quantities.

Byrne & Heavey (2006) show that sharing demand information in a 3-tier supply chain with restricted production capacity is reducing costs and increasing customer service level (CSL). Reduced inventory levels are also shown in similar settings by Cachon & Fisher (2000); Lau et al. (2004); Li et al. (2001); Yu et al. (2001), while the model of Li et al. (2001) highlights a possible negative impact of information sharing on CSL. Terwiesch et al. (2005) state that sharing of forecast information could increase SC performance in the semiconductor industry, but the potential remains unlocked due to the likelihood of silo-optimization of the supply chain partners. This highlights that information sharing mechanisms do not always establish a pareto-optimal situation for all SC partners.

Information technology can support the implementation of coordination mechanisms and increases the effectiveness and efficiency of information transfer between supply chain partners (Kanda & Deshmukh, 2008). A common application of the information technology measure is vendor managed inventory (VMI), which transfers inventory responsibility to an upstream part of the supply chain and results in cost savings

[1] According to Kanda & Deshmukh (2008), collaboration and coordination may be considered synonymous terms in this context, we will only use coordination as term in this paper.

through improved transparency about the inventory status along the SC (Haavik, 2000).

More advanced coordination methods include joint decision making between the supply chain partners. The objective shifts from optimizing a single entity to a holistic optimization of the whole SC. One method for joint decision making is collaborative planning, forecasting and replenishment (CPFR), which includes not only the sharing of demand information but also collaborative determination of forecasts and replenishment decisions (Skjoett-Larsen et al., 2003). Aviv (2001) examines the impact of sharing forecasts on service level and supply chain cost in a two-tier supply chain under uncertain demand. The resulting lower cost for all supply chain partners are complemented by improved agility of the manufacturer and improved forecast quality upstream of the SC. These findings are supported by Aviv (2007) as well as McCarthy & Golicic (2002). Chen & Chen (2005) demonstrate cost reductions through joint order cycle determination and quantity planning under varying demand, and Sari (2008) shows that using the same forecast for replenishment planning outperforms VMI. Kamalapur et al. (2013) show that especially in settings with high demand variation, limited production capacity, and high backorder cost, CPFR leads to a more cost efficient SC. Results from Sawik (2009) indicate that a chronological coordination of the planning process steps between the supply chain partners does not significantly reduce costs.

For the purpose of this research, the existing SC models lack important characteristics of today's networks to be found in the health care industry, such as considering more than two echelons (Kanda & Deshmukh, 2008). Coordination, taking into account the perishability of goods, has been covered to a limited extent only in the existing literature. Ketzenberg & Ferguson (2008) show that sharing inventory age information and VMI improve the effectiveness of a supply chain handling perishable grocery products through improved customer service level and reduced write-off costs; Karaesmen et al. (2011) highlight that supply chain coordination focusing on perishable goods should be a direction for future research.

Supply Chain optimization in the health care industry Several areas for improvement in the health care supply chain are discussed in the literature. Relevant examples are lack of application of collaborative supply chain management methods like CPFR (Aptel & Pourjalali, 2001), misaligned incentives across the supply chain

partners (McKone-Sweet et al., 2005), and insufficient inventory management leading to stock-outs or write-offs (Romero, 2013). Other examples focus on single parts of the supply chain like the inflexibility in manufacturing lead times and capacity (Ebel et al., 2013). Such improvement opportunities are required in an environment where operational efficiency is becoming increasingly important to be able to comply with tighter regulations, maintain quality of care, and avoid drug shortages (Alicke et al., 2014; Romero, 2013; Vila-Parrish et al., 2012).

The health care sector could benefit from as of today largely disregarded SC collaboration methods (de Vries et al., 2011), some of which have successfully been implemented in retail. Kiely (2004) states that the lack of coordination methods can be traced back to a relatively equal balance of power between the SC partners leading to a resistance to share more information than necessary.

Several authors have examined the applicability of lean supply chain methods in the health care sector. A case study conducted by Rivard-Royer et al. (2002) finds, with some limitations, that there is no major advantage of a zero inventory system based on pull replenishment in a hospital ward. Jarrett (1998) emphasizes the need to implement just-in-time in the health care supply chain to remain competitive.

McKone-Sweet et al. (2005) highlight misaligned incentives along the SC partners as one of the reasons for bad performance of the health care supply chains. Physicians tend to overstock to hedge against volatile demand, hospital inventory managers intend to drive down costs through reduced inventory levels, whereas manufacturers try to push as much inventory as possible into the supply chain to promote their products through increased availability.

According to Bhakoo et al. (2012), the level of coordination in the Australian health care supply chain is lower for the upstream parts of the chain. It is further influenced by the turnover and criticality of the product as well as the geographical complexity of the network and the mutual trust between the partners.

While potential benefits of SC coordination have been highlighted in the literature, evidence which conditions influence these benefits and how they are distributed among SC partners is sparse. We contribute to the literature through analyzing the impact of SC characteristics as demand variability, production capacity, or shelf life length, on the benefit of coordination schemes. We analyze changes in SC performance as

well as shifts in the cost distribution across the SC entities. To ensure the validity of our model, we use real-world demand data in a discrete simulation model which reflects the characteristics of a pharmaceutical SC.

2.3 Model development

Among methods available for supply chain modeling, discrete event simulation is a widely used method to analyze SC decision problems (Tako & Robinson, 2012). We choose a discrete event simulation approach, as an analytical solution would be intractable when accommodating the complex model structure required. Discrete event simulation has been used in several prior studies examining SC performance (Angulo et al., 2004; Fleisch & Tellkamp, 2005; Shin & Benton, 2004). In the remainder of this paper, the notation presented in Table 2.1 is used.

2.3.1 Structure of the Supply Chain under study

The modeled SC has a divergent, three-echelon structure (see Figure 2.1) and is tailored in its size and interdependencies to the network of German pharmacies and pharmaceutical wholesalers. Each pharmacy sources at one of the three wholesalers, which receive the corresponding pharmaceuticals from a manufacturer. Delivery frequencies can only be varied as integer multiples of the frequency of the next downstream echelon to maintain a consistent tact in the simulation model. This constraint poses no limitation to the application of the model, as delivery frequencies of the modeled SC increase with each upstream echelon. Production and transportation lead times are assumed to be constant over time but can vary between entities. Every period, each pharmacy is facing a share of the total demand which is proportional to its market share (see Section 2.3.4). Unsatisfied demand is lost at the pharmacies and wholesalers, as patients can buy the needed pharmaceuticals at a different pharmacy and pharmacies could source them at a different wholesaler. The manufacturer backlogs unsatisfied demand, as he is the single available supplier for a specific pharmaceutical. Pharmaceuticals have a fixed shelf life, ranging from 3 to 12 months. Each echelon can only deliver a pharmaceutical if the remaining shelf life is longer than the minimal remaining lead time for a possible patient delivery. Pharmaceuticals older than the

Table 2.1: Notation of variables used in the simulation model

Input variables		Model variables	
D_p	Demand in units at p	$B_{i,t}$	Backlog at i in t
$DE(i)$	Set of downstream entities connected to i including i	$CT^*_{i,t}$	Optimal time to cover with safety stock
DF_n	Delivery frequency of echelon n	F_{i,t,t_f}	Forecasted demand for i in t for t_f
$E(i)$	Connected entity i at next upstream echelon	$FE_{i,t,t_{fe}}$	Forecast error of i in t for a forecast of t_{fe} periods
HC_i	Holding costs per unit and period at i	FS	Forecast smoothing constant
LC_i	Logistic cost per unit at i	$IN_{i,t,a}$	Inventory in t at i with age a
LT_i	Lead time towards i	$LS_{i,t}$	Lost sales at i
$n(i)$	Echelon of entity i	MS_i	Market share of i
SL	Maximum shelf life	$O_{i,t}$	Order quantity of i
SOC_i	Stock-out cost per unit at i	$R_{i,t,a}$	Replenishment for i shipped at a
UC_i	Unit cost at entity i	$SS_{i,t}$	Safety stock at i in t
$UE(i)$	Set of upstream entities connected to i	$WIP_{i,t}$	Work in progress to i in t
α	Target Pharmacy service level for safety stock optimization	$WO_{t,i}$	Write-offs at i

Output variables		Indices	
FR_p	Fill rate of p	i	Supply chain entity, $i = 1, \ldots, I$ with $I = P + W + M$
HC_n	Holding costs at n	m	Manufacturer, $m = 1, \ldots, M$
IC_n	Inventory cost at n	p	Pharmacy, $p = 1, \ldots, P$
LC_n	Logistic costs at n	t	Simulation period, $t = 1, \ldots, T$
SCC_n	Supply Chain cost at n	w	Wholesaler, $w = 1, \ldots, W$
WOC_n	Write-off-costs at n		

shelf life are written-off.

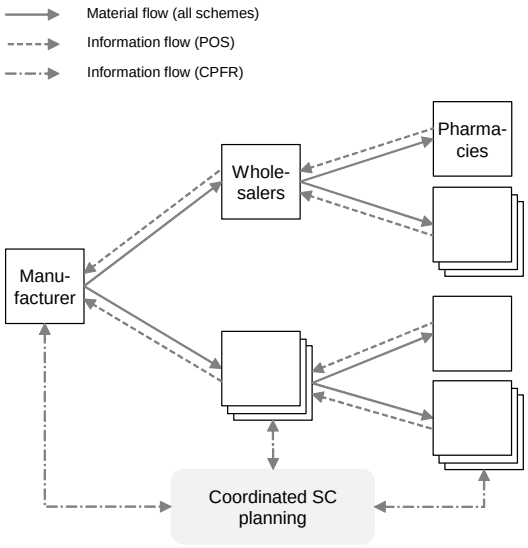

Figure 2.1: Structure of modeled supply chain

The discrete event simulation model simulates 281 decision events in one simulation period, representing 11 months of available demand data with 5 working days per week. The simulation runs three successive simulation periods to permit perishability, even when shelf life is longer than 11 months.

2.3.2 Simulation model

Model process The simulation is divided into four modules, each performing the outlined steps sequentially for every echelon:

1) **Write-off perished inventory.** Inventory, which's age is too high ($a \geq a_n^{max}$) to be sent from the echelon n to the patient with sufficient shelf life is written off, otherwise it is transferred to next period's inventory and aged on period:

$$IN_{i,t,a+1} = \begin{cases} IN_{i,t-1,a} & \forall\ a\ < a_{n(i)}^{max} \\ 0 & \forall\ a\ \geq a_{n(i)}^{max} \end{cases} \tag{2.1}$$

The maximum age a_n^{max} in each echelon is either the sum of the average remaining downstream lead times for the manufacturer (1) and wholesaler (2). As the

pharmacy lead time towards the patient is zero, $a_{n=3}^{max}$ is equal to the shelf life:

$$a_{n=1}^{max} = \overline{LT}_{i,i=(M+1,\ldots,M+W)} + \overline{LT}_{i,i=(M+W+1,\ldots,M+W+P)}, \tag{2.2}$$

$$a_{n=2}^{max} = \overline{LT}_{i,i=(M+W+1,\ldots,M+W+P)}, \tag{2.3}$$

$$a_{n=3}^{max} = SL. \tag{2.4}$$

2) **Receive replenishments.** Replenishments delivered to entity i, are added to the inventory in t when they have been shipped in period $t - LT_i$. Having been sent with the age a, they are added to the inventory having age $a + LT_i$:

$$IN_{i,t,a+LT_i} = IN_{i,t,a+LT_i} + R_{i,t-LT_i,a}. \tag{2.5}$$

3) **Satisfy pharmacy demand and ship manufacturer and wholesaler orders from last period.** Demand or orders are satisfied from the available inventory. To fulfill demand $D_{p,t,a}^{sat}$ pharmacies use their oldest inventory first. For each inventory age, either the maximum available inventory or the remaining unsatisfied demand is provided to the patient:

$$D_{p,t,a}^{sat} = \max\left(0, D_{p,t} - \sum_{b=a+1}^{SL} D_{p,t,b}^{sat} - \max\left(0, D_{p,t} - IN_{p,t,a} - \sum_{b=a+1}^{SL} D_{k,t,b}^{sat}\right)\right) \tag{2.6}$$
$$\forall\ a = (SL, \ldots, 1).$$

The remaining inventory is updated accordingly:

$$IN_{p,t,a} = IN_{p,t,a} - D_{p,t,a}^{satisfied}. \tag{2.7}$$

For the manufacturer and wholesaler echelons, respectively, three possible allocation policies are used if the inventory is not sufficient to provide replenishments for all outstanding orders. These include a simple first-come-first-serve (FCFS) policy, a proportional policy from Cachon & Lariviere (1999), and a market share based policy used in the German pharmaceutical market according to industry experts.

First-come-first-serve (FCFS). Orders are satisfied in the sequence of the entity

number i with inventory not older than $a^{max}_{n(i)}$:

$$R_{i,t,a} = \max\left(0, O_{i,t-1} - \sum_{b=a+1}^{a^{max}} R_{i,t,b} - \max\left(0, O_{i,t-1} - IN_{E(i),t,a} - \sum_{b=a+1}^{a^{max}} R_{i,t,b}\right)\right) \tag{2.8}$$

$$\forall\ a = (a^{max}, \ldots, 1).$$

Proportional allocation. For a more evenly distributed allocation of the available wholesaler and manufacturer inventory among their respective downstream echelons, we use a proportional allocation following Cachon & Lariviere (1999). Each order is only satisfied by entity i up to the proportion of the factor $\nu_i \in (0, 1)$, the share of all orders which can be satisfied through the inventory on hand with sufficient low age $a \leq a^{max}$:

$$\nu_i = \min\left\{1, \frac{\sum_{a=1}^{a^{max}} IN_{i,t,a}}{\sum_{j} O_{j,t-1}}\right\}\ \forall\ j \text{ with } E(j) = i \text{ and } n(i) = 1, 2. \tag{2.9}$$

Under the proportional allocation policy, the outstanding orders are realized as under the FCFS policy, with the order quantity corrected by ν_i is shown in Equation (2.10):

$$R_{i,t,a} = \max\left(0, O_{i,t-1}\nu_i - \sum_{b=a+1}^{a^{max}} R_{i,t,b} - \max\left(0, O_{i,t-1}\nu_i - IN_{E(i),t,a} - \sum_{b=a+1}^{a^{max}} R_{i,t,b}\right)\right). \tag{2.10}$$

Market share based allocation. A third allocation policy is based on interviews with experts from German pharmaceutical wholesalers. Each entity can only order a share of the available inventory equivalent to its market share and no more than 150 % of the mean of the past ρ orders:

$$O^{adjusted}_{i,t-1} = min\left\{O_{i,t-1}, 1.5 \cdot \frac{\sum_{t'=t-\rho-1}^{t-1} O_{i,t'}}{\rho}, MS_i IN_{E(i),t,a}\right\}. \tag{2.11}$$

This policy is enforced through setting $O_{i,t} = O^{adjusted}_{i,t}\ \forall\ n(i) = 1, 2$ in Equation (2.8).

For each of the three aforementioned allocation policies, the inventory of the replenishing entities is updated accordingly:

$$IN_{i,t,a} = IN_{i,t,a} - \sum_{j} R_{j,t,a} \ \forall \ j \text{ with } E(j) = i. \tag{2.12}$$

4) **Define order quantity.** The order quantity is computed through three different schemes with increasing SC coordination.

Scheme 1 - conventional supply chain information sharing (CONV). The order quantity is defined based on an adaption of the extended newsvendor approach from Kanchanasuntorn & Techanitisawad (2006) with:

$$F(Z_i) = \frac{SOC_i}{SOC_i + WOC_i + HC_i}. \tag{2.13}$$

$F(Z_i)$ is the critical fractile, Z_i is be obtained from the inverse cumulative distribution function of the demand distribution. We fit log-normal distributions to the demand patterns of each product (see Section 2.3.4). The ordering decision for all echelons are characterized by Equation (2.14), where the stock shipped from the upper echelon as well as the available inventory is subtracted from the order-up-to level:

$$\begin{aligned} O_{i,t} = \max(0, \mu_{D_{i,t}}(DF_n + LT_i) + \sigma_{D,i,t}\sqrt{(DF_n + LT_i)}Z_i \\ - WIP_{i,t} - \sum_{a=1}^{SL} IN_{i,t}). \end{aligned} \tag{2.14}$$

Scheme 2 - point of sale information sharing (POS) The order quantity calculation uses a forecast which is based on the demand data collected by the pharmacies. Considering seasonal characteristics of some of the pharmaceuticals modeled (e.g. allergy or cold medicine), a Holt-Winters forecasting method is used. For a comprehensive description of the Holt-Winter method, we refer to Holt (2004).

We use a Holt-Winters algorithm with additive factors to separate the level, trend, and seasonal components of the time series. The constants $A, B, C \in [0, 1]$ are smoothing factors for the level, trend and seasonal change of the time series, respectively. They are set as follows: $A = 0.5$, $B = 0$, and $C = 0.8$, which results

into a reliable replication of demand curves tested during our experiments. B is set to zero as we do not observe any trend effect in our data. The seasonal cycle is synchronized with the simulation period of 281 days. The update equations of the Holt-Winters forecast are initialized with zeros and the forecast is run for six periods for calibration prior to our experiments. The order quantity is then determined using the forecasted demand quantity and an additional safety stock.

$$O_{i,t} = \max\left(0, \sum_{t_f=t+1}^{t+LT_i+DF_n} F_{i,t,t_f} + SS_{i,t} - WIP_{i,t} - \sum_{a=1}^{SL} IN_{i,t}\right). \tag{2.15}$$

The safety stock quantity is computed through the product of the forecast error's standard deviation over the lead time σ_{FE} and the safety stock factor SSF (Silver & Peterson, 1985). F^{-1} is the inverse of the cumulative distribution function of the demand distribution:

$$SS_{i,t} = \sigma_{FE} SSF \text{ with } SSF = F^{-1}\left(\frac{HC_i}{SOC_i}\right). \tag{2.16}$$

Scheme 3 - collaborative planning, forecasting and replenishment (CPFR). The calculation of the order quantity incorporates the forecasted demand of the pharmacies connected over the lead time to SC entity i, the inventory and work in progress of all connected downstream entities, as shown in Equation (2.17):

$$\begin{aligned} O_{i,t} = \max\Bigg(0, \sum_{j \in DE(i)|n(i)=3} \sum_{t_f=t+1}^{t+LT_i+DF_{n(i)}} F_{j,t,t_f} + SS_{i,t} \\ - \sum_{j \in DE(i)} \sum_{a} IN_{j,t,a} - \sum_{j \in DE(i)} WIP_{j,t}\Bigg). \end{aligned} \tag{2.17}$$

The safety stock is allocated across entities using a Dynamic Programming (DP) algorithm based on Minner (1997); Inderfurth & Minner (1998). To limit the size of the action space, the algorithm utilizes a finding from Simpson (1958). He shows that the lead time of the upstream echelons, or the time from order to replenishment, $CT^*_{i,t}$, against which the safety stock covers, should either be 0 or the sum of all upstream lead times. Thus, every pharmacy covers its lead time plus the sum of all uncovered lead times of its assigned upstream echelon entities. The optimal coverage time for a wholesaler or manufacturer is therefore either zero or the sum of all uncovered lead times of the assigned upstream echelon

entities including the own lead time. In our model, we derive $CT^*_{i,t}$ as described in Minner (1997) and refer to his article for a detailed description of the DP algorithm.

The above described modules 1-4 are executed sequentially in every simulation period $t = 1, \ldots, T$.

2.3.3 Performance measures

Two performance measures are calculated, supply chain costs (SCC) and fill rate (FR). SCC are be calculated over all echelons, while FR is only calculated for the patient demand at the pharmacy echelon. Holding cost are charged for each item on stock at the beginning of a simulation period, the sum over all entities and all simulation periods results in the total inventory cost IC. Logistic cost LC are incurred when a replenishment is shipped between SC entities. The write-off cost WOC are the sum of the product of written off pharmaceuticals in each entity with the respective write-off cost. The sum of the three cost measures results in the total SC cost SCC.

The SCC are calculated as shown in Equations (2.18 - 2.22) and the FR is calculated through Equation (2.22).

$$IC = \sum_t \sum_i \left(\sum_a IN_{i,t,a} \right) HC_i \tag{2.18}$$

$$LC = \sum_t \sum_i \left(\sum_a R_{i,t,a} \right) TC_i \tag{2.19}$$

$$WOC = \sum_t \sum_i WO_{i,t} UC_i \tag{2.20}$$

$$SCC = IC_n + LC_n + WOC_n \tag{2.21}$$

$$FR = 1 - \sum_t \sum_p \left(\frac{LSp,t}{D_{p,t}} \right) \tag{2.22}$$

$$\text{with } p = (1, \ldots, P),\ i = (1, \ldots, I) \text{ and } t = (1, \ldots, T)$$

The fill rate FR is 1 minus the share of unsatisfied demand, summed over all periods and pharmacies.

2.3.4 Simulation input

Demand data To ensure the validity of the model inputs, real demand data from the German pharmaceutical market is used (data has been provided by the companies *Insight Health* and *ABDATA*, see Appendix A.1 for a detailed description of the dataset). The selected dataset contains monthly sales quantities over 11 months for three different pharmaceutical products in 16 German states. Figure 2.2 shows a plot of the timeline for the pharmaceutical with pointed demand, a product used by patients suffering from allergies.

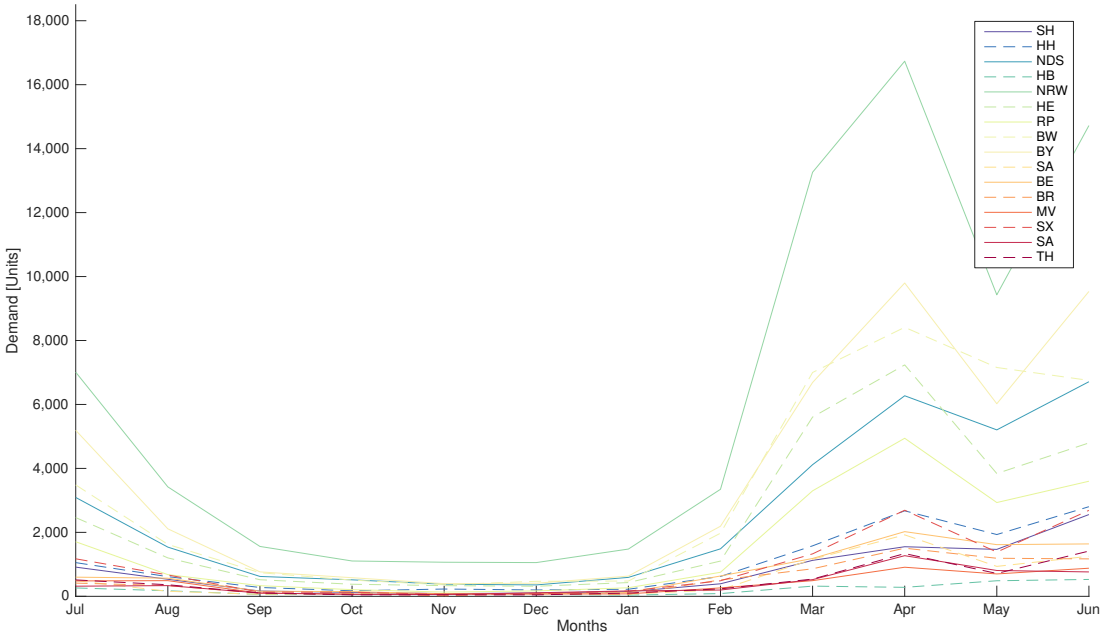

Figure 2.2: Demand pattern of pharmaceutical with pointed demand for the 16 German states[2].

To approximate the sales quantity for a single pharmacy, the aggregated state demand is distributed to the pharmacies in each state following the revenue distribution of the German pharmacies (ABDA, 2013). The pharmacy revenue data is fitted with a log-normal distribution. The goodness-of-fit is tested through a Kolmogorov-Smirnov test, which confirms the fit of the distribution at a 1% significance level. The data update frequency is increased from monthly to daily updates through cubic interpolation. The absolute demand quantity is normalized across pharmaceuticals to ensure comparability of the resulting SCC.

[2] The plots for the remaining pharmaceuticals can be found in Appendix A.1.

To produce i.i.d. random samples of the demand time series to be used in the different simulation iterations, a maximum entropy bootstrap procedure was used, as it is able to bootstrap non-stationary data. Vinod (2004) defines the bootstrap algorithm in seven steps: (1) Arrange the data in an increasing order, (2) determine the intermediary points between the arranged data, (3) calculate the truncated mean across all observations and use it to compute the lower and upper intermediary points, (4) calculate the average of the maximum entropy density for each interval, (5) use a uniform random distribution $U[0, 1]$ to calculate sample quartiles for the maximum entropy density and sort them, (6) sort the quantiles re-establishing the order of the original dataset, (7) repeat step one to six for the number of required bootstraps. Bootstrap samples of the data for pharmaceutical 3 are shown in Figure 2.3. In our experiments, this method produced superior time series bootstrapping results compared to approaches using autoregressive functions.

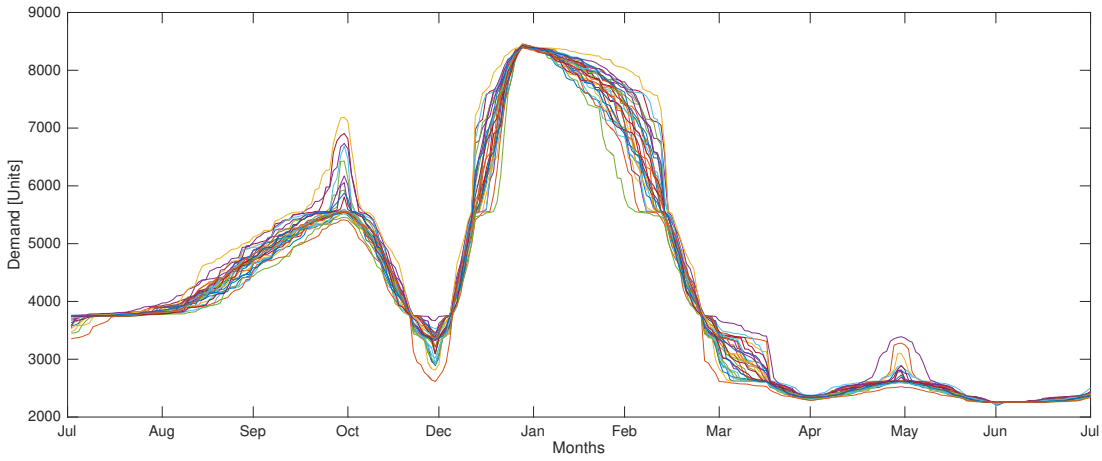

Figure 2.3: 30 samples of bootstrapped demand data of the pointed pharmaceutical, state RP (Rhineland-Palatinate) using the maximum entropy boostrap method.

Supply Chain parameters SC parameters from the German pharmaceutical market are used in the model. They have been obtained through interviews with experts from pharmacies as well as pharmaceutical wholesalers and are presented in Table 2.2. Three out of 16 German states are modeled (Hamburg, Mecklenburg-Western Pomerania, and Rhineland-Palatinate), representing the spatial heterogeneity population and pharmacy density among German states. The number of pharmacies is reduced by factor 4 to make the simulation computational feasible and to account for the fact that not all pharmacies are using all of the 13 wholesalers operating on the

German market (PHAGRO, 2016). Unit cost are assumed to be equal for all product to ensure comparability of the resulting SCC.

Table 2.2: Parameter values for three entities in the simulation

Entity	Parameter	Value	Source
Pharmacies	# of Pharmacies	477	ABDA (2013)
	LT from Wholes. [days]	1	Expert Interview
	Delivery frequency [days]	1	
	HC [share of unit cost]	10%	
	TC [EUR per unit]	0.16	
	UC [EUR]	10.21	*ABDATA* dataset[1]
	CLS [EUR]	777	
Wholesalers	# of Wholesalers	3	Wholesaler[2]
	LT from Manufact. [days]	7	Expert Interview
	Delivery frequency [days]	8	
	HC [share of unit cost]	7%	
	TC [EUR per unit]	0.4	
	UC [EUR]	8.65	*ABDATA* dataset
Manufacturer	# of Manufacturers	1	Expert Interview
	Manufacturing LT [days]	30	
	Production frequency	16	
	HC [share of unit cost]	10%	
	UC [EUR]	1.73[3]	

[1] Cost of lost sales are calculated through subtracting the unit cost from the pharmacie's selling price
[2] www.phoenixgroup.eu, accessed October 30th, 2015
[3] Assuming an 80% margin of the manufacturer, logistic cost included

2.4 Design of experiments

Based on the results of the literature review, six hypotheses are proposed about the response of the two dependent variables SCC and FR with five independent variables ALL, CAP, P, SCS and SL:

As SC coordination in general has been identified as cost-effective we follow Byrne & Heavey (2006); Cachon & Fisher (2000); Lau et al. (2004); Terwiesch et al. (2005) by proposing hypothesis 1:

H1a: CPFR entails the lowest SCC followed by POS and CONV.

H1b: CPFR entails the highest FR followed by POS and CONV.

Sari (2008); Kamalapur et al. (2013) state that the variability of demand is handled better in coordinated SCs. We use that finding for our second set of hypotheses:

H2a: Higher demand variability increases SCC to a lesser extent for CPFR than for POS and CONV.

H2b: Higher demand variability reduces FR to a lesser extent for CPFR than for POS and CONV.

Ketzenberg & Ferguson (2008) show that the cost-effectiveness of sharing inventory age information, especially for products with a short shelf life.

H3a: Lower SL deteriorates SCC to a lesser extent for CPFR than for POS and CONV.

H3b: Lower SL deteriorates FR to a lesser extent for CPFR than for POS and CONV.

Sari (2008) states that benefits of coordination schemes are highest when sufficient capacity is available. We follow her and test whether coordination schemes provide better mitigation of this effect through coordination:

H4: The benefits of CPFR and POS for the FR decrease when production capacity is reduced.

The proportional and practical allocation methods distribute the products more equally across the entities of an echelon. Due to the increased average availability, we propose a positive effect of the proportional and practical allocation techniques on overall fill rate as compared to FCFS allocation:

H5: Using proportional allocation (PROP) or practical allocation (PRAC) increases the FR compared to a FCFS allocation.

Chiang & Feng (2007) show that information sharing is more beneficial for the entities upstream of the supply chain. We therefore propose that the share of the benefits through using CPFR is higher for entities further upstream of the supply chain.

H6: The higher the position upstream in the SC, the higher the share of the coordination benefit.

To test these hypotheses, five independent variables are used in the main experimental design: Type of supply chain scheme (SCS), product (PROD), capactiy (CAP), shelf life (SL), and allocation (ALL). SCS indicates the coordination scheme under which the supply chain is operating: conventional supply chain information sharing (CONV), point of sale information sharing (POS), or collaborative planning, forecasting and replenishment (CPFR). PROD selects one out of 3 pharmaceutical products, indicating the demand pattern. CAP regulates the available capacity relative to the demand. The length of the shelf life is determined by the variable SL. If capacity is allocated through FCFS, proportional or market-share allocation is set by the variable ALL. Table 2.3 shows the treatment levels of the independent variables. The two dependent variables are SCC, summing up the cost of the whole supply chain and FR, indicating the percentage of patient demand covered in the period of the demand.

To evaluate H6, we test whether the cost share of the manufacturer (CSM) and pharmacies (CSPH) changes when using different coordination schemes.

Table 2.3: Independent variables and respective treatment levels included in the experimental design.

Independent variables	Treatment levels		
	1	2	3
PROD	1	2	3
SL [days]	281	140	70
CAP [Demand]	0.9	1.1	1.3
ALL	FCFS	PRAC	PROP
SCS	CONV	POS	CPFR

2.5 Model output analysis

To test the significance of the effects in the simulation results, we use a Multivariate Analysis of Variance (MANOVA), a recommended tool for simulation output analysis (Balci, 1990). It controls the type I error rate when using multiple dependent variables and can be applied to factorial experimental designs (see Hair et al., 2006 for

a detailed methodological description). When using MANOVA, four main assumptions have to be considered according to Hair et al. (2006): (1) Observations should be independent. As we use i.i.d. bootstraps in our experiments, we can assume the independence of observations in the simulation output. (2) Correlation of dependent variables. Following Hair et al. (2006) we use Bartlett's test for sphericity to assess this assumption. It tests the dataset for correlations and significant intercorrelation among the dependent variables. Results indicate an intercorrelation among the dependent variables with a significance of .000, with the correlation between SCC and FR being 0.52. The usage of MANOVA is acceptable when the dependent variables are moderately intercorrelated. (3) Multivariate normality of dataset. As there is no direct test available, we test the univariate normality of the results in each treatment group using a Jarque Bera test, giving an indication for the compliance to this assumption. Results show that the assumption of normality is not met by 29% of the treatment groups in our dataset at the 5% level. As the overall sample size is large, we are still able to perform MANOVA. (4) Homogeneity of the dependent variable's variance-covariance matrices between treatment groups. To test this assumption we would normally use the Box's M test, but it is highly sensitive to violations of the normality assumptions, see (3). We therefore omit the test due to the fact that we have equal treatment group sizes and a violation of this assumption would have minimal impact on the results. As MANOVA is very sensitive to outliers, we truncate our dataset beyond the 5% and 95% percentile.

We run $N = 30$ iterations of each of the $G = 243$ treatment groups using Matlab, R2014a, as the computation engine (see Figure 2.4 for the process of the experiment). For analyzing the output of the simulation, the MANOVA algorithms provided by STATA, Version 14, were applied. We use one MANOVA including the two dependent variables to test our hypotheses applying a significance level of 0.01 to test the significance of group differences[3]. To increase validity of our results we also perform univariate analysis of variances (ANOVAs) on each dependent variable, supporting the validity of our results (see Section A.2.2 for the detailed ANOVA results).

The MANOVA results in Table 2.4 show that the group differences relevant for the constructed hypotheses are significant at the required level in all four relevant tests (Wilks' lambda, Pillai's trace, Lawley-Hotelling trace, and Roy's largest root). This

[3] Although the non-compliance with three out of four assumptions is acceptable due to our experimental design, we chose the strict significance level 0.01 to avoid accepting nonsignificant differences.

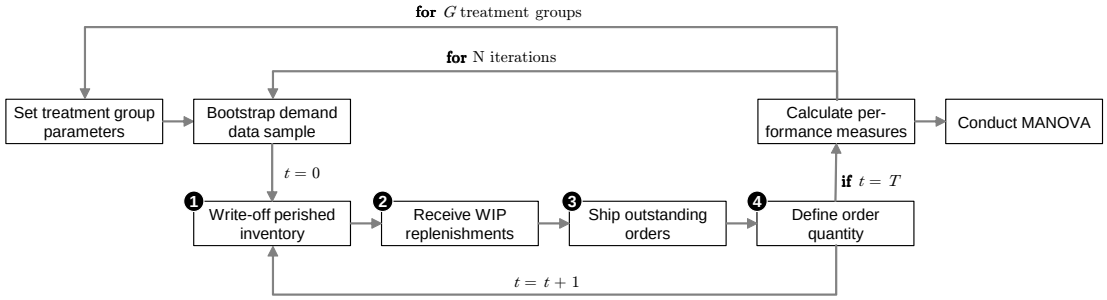

Figure 2.4: Modules and flow of the simulation model and design of experiment.

finding is supported by the results of the ANOVAs (Table A.7 - A.10). This implies that the choice of coordination scheme for the SC has significant impact on the cost and fill rate, respectively. Further, the resulting cost difference depends on the demand variability, product shelf life, production capacity and allocation method. The following paragraphs present the results for each hypothesis. Exact estimates of the dependent variables are presented in Section A.2.1, supporting the results of Fig. 2.5 - 2.9.

H1 - Impact of the degree of coordination Model results in Figure 2.5 show that coordination through POS and CPFR achieves lower SCC and a better FR compared to the CONV scheme without coordination. We can therefore confirm H1, with our findings being in line with Byrne & Heavey (2006); Cachon & Fisher (2000); Lau et al. (2004); Terwiesch et al. (2005). The cost benefit of POS is accomplished through a shift of inventory from the two upper to the pharmacy echelon, resulting in 59% lower total inventory cost. This leads to 20% higher logistic, but also 84% lower write-off cost, resulting in an overall cost advantage of the POS against the CONV scheme of 36%. The CPFR scheme's benefit with 40% is even larger. Compared to the POS scheme inventory cost are reduced by 20%, with substantial reduction at the wholesaler and pharmacy echelon and an increase at the manufacturer echelon. The second significant cost benefit against the POS scheme are the lower write-off cost, particularly at the wholesaler and manufacturer echelon. The intensified coordination allows an increase of the fill rate from 67% (CONV) to 87% (POS) and 90% (CPFR). Based on these results, we conclude that both coordination schemes provide overall value for a pharma supply chain. The numerical results for H1 are presented in Table A.1.

Table 2.4: MANOVA results for the main experimental design showing the significance of the effects necessary for testing hypotheses 1 – 6

Effect	Statistics		F value	Prob > F
SCM	W	0.103957	3662.415	.0000 [e]
	P	0.932322	1522.031	.0000 [a]
	L	8.27041	7205.594	.0000 [a]
	R	8.227996	14341.40	.0000 [u]
All	W	0.873195	122.2526	.0000 [e]
	P	0.127391	118.5736	.0000 [a]
	L	0.144549	125.938	.0000 [a]
	R	0.139748	243.5814	.0000 [u]
PRO × SCM	W	0.326123	383.8127	.0000 [a]
	P	0.7282	258.6988	.0000 [a]
	L	1.900701	551.9556	.0000 [a]
	R	1.810004	2103.829	.0000 [u]
CAP × SCM	W	0.266797	466.586	.0000 [a]
	P	0.887333	331.3484	.0000 [a]
	L	2.186003	634.8061	.0000 [a]
	R	1.903016	2211.939	.0000 [u]
SL × SCM	W	0.867087	63.61098	.0000 [a]
	P	0.136314	61.51192	.0000 [a]
	L	0.149375	65.06701	.0000 [a]
	R	0.116759	203.5699	.0000 [u]
ALL × SCM	W	0.675355	122.2349	.0000 [a]
	P	0.349592	112.7509	.0000 [a]
	L	0.444184	130.6549	.0000 [a]
	R	0.342302	403.0036	.0000 [u]

[e] exact
[a] approximate
[u] upper bound on F
W = Wilks' lambda
P = Pillai's trace
L = Lawley-Hotelling trace
R = Roy's largest root

H2 - Ability of SC coordination to cope with variable demand Results in Figure 2.6 present the ability of the different coordination schemes to handle fluctuating demand. The cost for the CONV scheme are steadily increasing as demand is becoming more variable (13% from the stable to the pointed pharmaceutical). This cost increase is mainly due to an overall higher inventory level (67% more inventory for the pointed compared with the stable pharmaceutical), especially at the manufacturer echelon (plus 298%), to buffer against the fluctuating demand. While this reaction

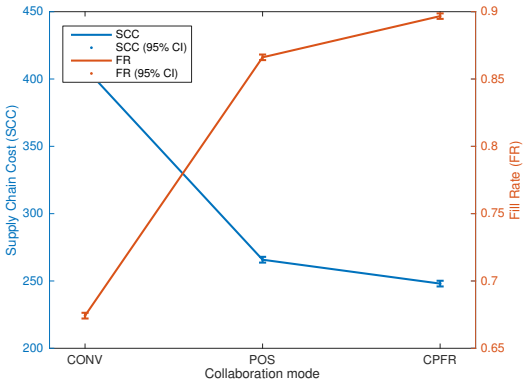

Figure 2.5: Estimated average SCC and FR using the three coordination schemes

also leads to higher write-off cost (plus 56%), it prevents a further decrease of the FR when variability increases from the seasonal to the pointed pharmaceutical. The cost difference between CPFR and CONV increases from 37% to 44% when going from the stable to the pointed pharmaceutical. This difference provides evidence for a higher value of a coordination scheme in markets with variable demands. The lower cost of the CPFR for the demands with medium and high variability (seasonal and pointed pharmaceuticals) compared to POS let us confirm H2a, in line with Sari (2008).

The FR decreases monotonically for both, the POS and the CPFR schemes. While the difference between the two schemes is 7 percentage points for the seasonal pharmaceutical, it is reduced to near zero for the pointed pharmaceutical. Both schemes deliver a higher FR compared to the CONV scheme, with the difference of the CONV to the CPFR scheme being highest for the seasonal pharmaceutical (34 percentage points). Results therefore indicate that increased coordination is not protecting the supply chain against a deteriorating FR when facing more variable demand, but the FR remains on a higher level for both schemes with coordination. We therefore have to reject H2b, but conclude that it would be beneficial to use coordination as well in markets with variable demand as it still provides a higher fill rate. Numerical results for H2 are presented in Table A.2.

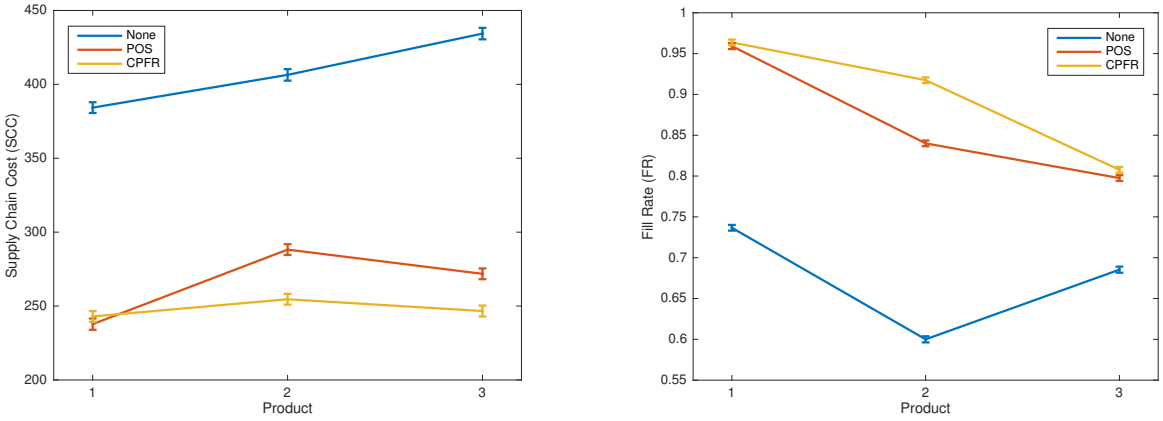

Figure 2.6: Estimated SCC (left) and FR (right) for the three coordination schemes and products with different demand variability

H3 - Influence of shelf life length on the benefits of SC coordination Results in Figure 2.7 show that the cost benefit of the schemes with coordination increases with shorter shelf life and they are also able to maintain a better FR across the shelf lifes tested, although the advantage decreases when shelf life is shortened. With the CONV scheme costs increase by 69% when shelf life is reduced from 12 to 3 months, mainly due to increased write-off cost (plus 102%). For the POS and CPFR scheme costs increase only by 6% and 5%, respectively. The FR remains stable across all three shelf lifes for POS and CPFR, while it increases by 11% for CONV when shelf life is reduced from 12 to 3 months due to larger inventory quantities at the manufacturer and pharmacy echelons. With the FR effect sizes for POS and CPFR being as small as 1%, we cannot draw conclusions beyond the observation that coordination schemes retain their advantage of a higher FR, even for products with shorter shelf life. We therefore can accept H3a but have to reject H3b. Considering the cost advantage of the schemes applying coordination we conclude that it is beneficial to apply coordination when handling products with a limited shelf life. Our findings are in line with Ketzenberg & Ferguson (2008), who also find that a central coordination scheme is most beneficial when handling perishable goods. Numerical results for H3 are presented in Table A.3.

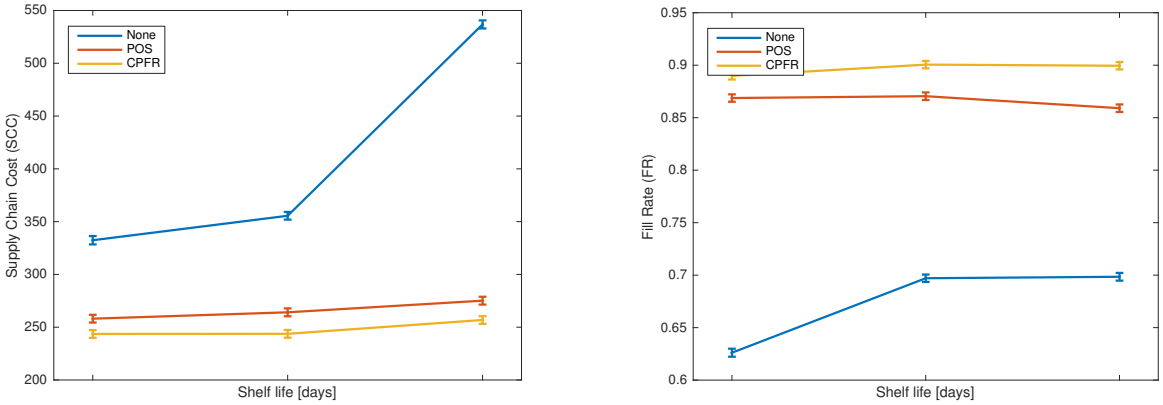

Figure 2.7: Estimated SCC (left) and FR (right) for the three coordination schemes and different shelf life lengths

H4 - Handling capacity constraints through SC coordination Results in Figure 2.8 show that the schemes using coordination dominate the CONV scheme for all production capacity levels tested. With limited manufacturing capacity, the fill rate subsequently declines for all schemes tested. The rate of decline increases with the degree of coordination. When reducing capacity from 1.3 to 0.9 × mean demand, the FR is reduced by 4, 10, and 13 percentage points for the CONV, POS, and CPFR scheme. With these results we can confirm H4b. When acting in an environment which requires a high and stable FR, like the pharmaceutical supply chain, implementation of coordination schemes would still be beneficial in environments with frequent manufacturing constraints, but a trade-off with the implementation cost for a coordination system (e.g. telecommunication infrastructure, additional headcount) has to be made. Detailed numerical results for H4 can be found in Table A.4.

H5 - Impact of allocation methods on SC performance Results in Figure 2.9 (left) show that the practical allocation performs worst among all allocation methods, with the difference being 6.4 percentage points compared to the FCFS method used in the POS scheme. Proportional allocation has a positive effect on the fill rate for the CONV coordination scheme (plus 3 percentage points), but for the two schemes with coordination the size of the effect is not big enough to allow a well-founded interpretation. We therefore cannot conclude that when using a coordination scheme in the supply chain, a more advanced allocation scheme would provide any significant value-add to performance achieved with the coordination, leading to a rejection of

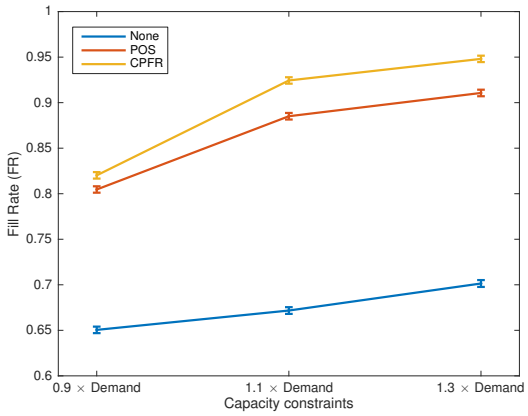

Figure 2.8: Estimated FR for the three coordination schemes and different production capacity constraints

H5. When coordination is not used, a proportional allocation method could provide value, as long as the behavior of the supply chain partners is focused on increasing the value for the whole supply chain and they not only make their decision based on the outcome for their own entity. Detailed numerical results are presented in Table A.5.

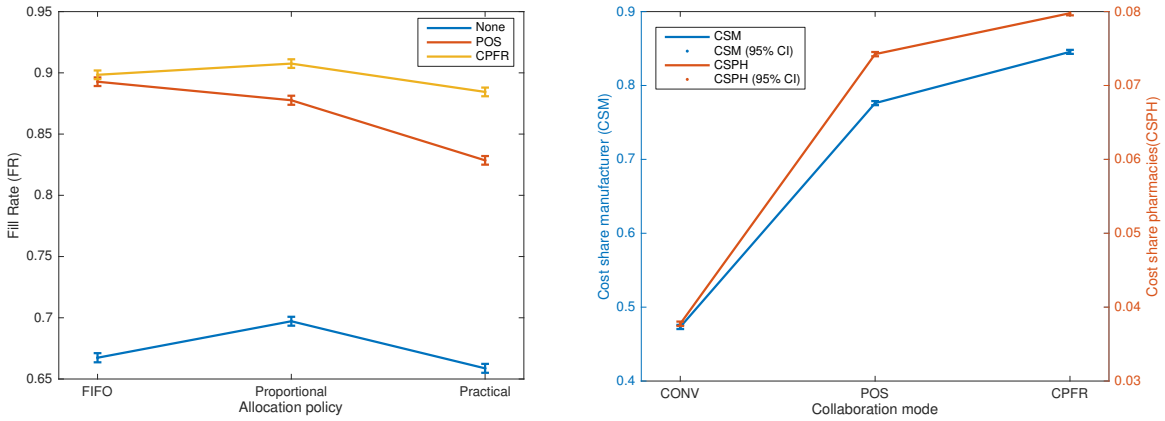

Figure 2.9: Left: Estimated manufacturer and pharmacy share of the total supply chain cost; Right: Estimated FR for the three allocation policies and coordination schemes

H6 - Impact of SC coordination schemes on cost distribution among SC partners Results in Figure 2.9 (right) show that coordination shifts costs back from the wholesaler to the upper echelons of the supply chain. Without coordination, cost are concentrated among the manufacturer (43%) and wholesalers (53%), with the pharmacies incurring only 4% of the cost due to substantially lower holding and

logistic cost. The increase of manufacturer and pharmacy cost share when using POS or CPFR reduces the wholesaler's cost share by up to 35 percentage points. These results do reject H6, as the only echelon with an absolute benefit from POS or CPFR is the wholesaler. With CPFR, the wholesaler's costs decrease by 91% while the cost of the manufacturer and pharmacy echelon increase by 15% and 35%, respectively.

2.6 Conclusion

Our study evaluates performance differences between three supply chain coordination schemes, with an increasing degree of coordination. The performance of each scheme is tested in a study simulating the characteristics of the German pharmaceutical market with varying demand variability, product shelf life, and production capacity. The main findings are:

The most-advanced coordination scheme, CPFR, is dominating the other two schemes by having the lowest cost and highest fill rate in almost all tested treatment groups. Leaving implementation aside, CPFR's benefits are clearly visible in the results.

The benefits of CPFR and POS against the CONV scheme vary when the independent variables are changed. The usage of POS or CPFR is highly advisable for products with a high variation in demand and a short shelf life, as both schemes keep cost low under these circumstances. If production capacity is scarce, coordination schemes should not be used as mitigation measure, as they perform best when ample capacity is available.

The reactions of the two schemes with coordination on changing conditions are comparable, with CPFR having an advantage of up to 12% (SCC) or 7.6 percentage points (FR) above POS, which already captures on average 89% of the benefit obtained through CPFR. The additional benefit of CPFR requires implementation of a complex management system across different companies, which generates additional cost for technical implementation (e.g. synchronizing data exchange), reduces control over proprietary information (e.g. POS data), and generates a need for cross-functional collaboration in the participating companies (Fliedner, 2003 and Skjoett-Larsen et al., 2003). A trade-off has to be made, whether this additional performance gain justifies such efforts.

With the implementation of coordination schemes, supply chain cost are shifted from the wholesaler to the manufacturer and pharmacy echelons. To motivate these two echelons to implement a coordination scheme, a benefit sharing agreement has to be reached, otherwise the pharmacies and manufacturer will not relinquish their dominant market positions.

The developed simulation model incorporates many characteristics of a real supply chain, but is also subject to limitations. The simulated echelons do only consider only one manufacturer and a limited set of wholesalers and pharmacies. To assess a potential effect of coordination in a competitive market, competing companies would need to be incorporated in the model. We ignored limitations regarding the timing of production and variability of transportation times to limit the complexity of the model. Incorporating these effects would further increase the robustness ot the results. To be able to test solely the effect of coordination, we neglected any implementation cost of the coordination schemes. These cost are a part of the overall SC cost and have to be considered when deciding about a SC scheme.

This work provides several opportunities for further research. A possible model extension could incorporate effects of benefit sharing contracts or evaluate the effect of a partial implementation of coordination across the whole SC. The performance of the tested coordination schemes could also be tested in different markets where perishability is of high relevance, e.g. consumer retail. Incorporating other factors influencing SC performance through uncertainty could increase the robustness of our findings. Variable lead times would induce uncertainty from the upper echelons of the SC and a more unstable manufacturing process would reflect the complexity of producing pharmaceuticals. Another possible extension could include the assessment of the change in effect size if only a share of the SC entities take part in the coordination effort, while others stick with the conventional SC structure.

Chapter 3

Was Angelina right?
Optimizing cancer prevention strategies for BRCA carriers

with Elisa F. Long and Stefan Spinler

The content presented in this chapter is based on Nohdurft et al. (2016b).

Female carriers of a BRCA1/2 genetic mutation face significantly elevated risks of cancer, with 45-65% of women developing breast cancer and 15-39% developing ovarian cancer in their lifetimes. Cancer risk can be reduced through prophylactic surgeries like a bilateral mastectomy (BM), bilateral salpingo-oophorectomy (BSO), or both. We develop a Markov Decision process (MDP) which models a carriers health states and uses clinical data to compute the optimal surgery sequence that maximizes the carrier's expected lifetime QALYs under varying assumptions about each surgery's impact on quality-of-life. Baseline results demonstrate that a QALY-maximizing sequence recommends BM between ages 30 and 60 and BSO after age 40. Surgeries are recommended later for BRCA2 carriers, as their risk for both cancers is less than for BRCA1 carriers. We derive structural properties from the model and show that when one surgery has already been completed, there exists an optimal control limit beyond which performing the other surgery is always QALY-maximizing.

3.1 Introduction

One out of every 400 women in the U.S.—and one in 50 women of Ashkenazi Jewish descent—carries a BRCA1/2 genetic mutation (Hall et al., 2009). Female carriers face a significantly higher risk of breast and ovarian cancer relative to the general population. An estimated 55-65% of BRCA1 mutation carriers and 45% of BRCA2 mutation carriers develop breast cancer in their lifetime, compared to 12% in the general population (Antoniou et al., 2003; Chen & Parmigiani, 2007). The lifetime risk of ovarian cancer is slightly lower at 39% (BRCA1) and 11-17% (BRCA2), but far greater than the 1-2% risk in the general population (Antoniou et al., 2003; Chen & Parmigiani, 2007).

Women with a known BRCA mutation can improve early cancer detection by undergoing enhanced surveillance with frequent mammograms and MRIs, and potentially reduce the risk of developing cancer with chemoprevention and/or prophylactic surgery. Surgically removing both breasts, known as a bilateral mastectomy (BM), can reduce a carrier's lifetime risk of breast cancer by as much as 95%; removing both ovaries and fallopian tubes, known as a bilateral salpingo-oophorectomy (BSO), can reduce ovarian cancer risk by 80% and additionally reduce breast cancer risk by up to 60% (Domchek et al., 2006, 2010; Eisen et al., 2005; Grann et al., 1999a; Rebbeck et al., 2009). As both surgeries are invasive procedures with possible complications, hormonal side effects, and fertility implications, deciding whether to undergo such procedures—and at what age—is a difficult choice for many women. The benefits of decreased cancer risk and peace of mind are often weighed against a reduction in quality-of-life (QOL). BRCA mutation carriers will increasingly face this trade-off, with more women learning of their own mutation status before a cancer diagnosis as less expensive genetic testing becomes widely available (Long & Ganz, 2015).

While the risk-reducing benefits of BM and BSO are widely accepted, no detailed guidelines currently exist to advise women at which age each surgery should be performed. The National Comprehensive Cancer Network (NCCN) advises physicians to discuss the option of BM with BRCA carriers, and suggests that a BSO be performed between ages 35 and 40 or after completion of child-bearing (National Comprehensive Cancer Network, 2016). Yet these recommendations do not differentiate between BRCA1 and BRCA2, specify a recommended age to undergo BM, nor incorporate patient QOL considerations. Other publications offer comparisons of cancer preven-

tion strategies (Schrag et al., 1997; Grann et al., 1999b), or provide information about cancer and mortality risk under different strategies (Kurian et al., 2014). However, a comprehensive model to optimize patient outcomes such as expected quality-adjusted life years (QALYs) or survival probability is still lacking.

In this study, we contribute to the existing literature by developing a comprehensive Markov Decision Process (MDP) model of a BRCA1 or BRCA2 mutation carrier's health states, and the potential impact of decisions to undergo prophylactic surgery on both breast and ovarian cancer risk, survival, and QOL. As the problem's state space is too large to solve directly with DP, we transform the problem of optimizing health outcomes into a form solvable through linear programming (LP) and we identify a QALY-maximizing sequence of prophylactic surgeries. To account for the uncertainty about the numerically estimated model parameters and carrier-specific preferences of the impact of surgeries on QOL, we conduct various sensitivity and robustness analyses. We additionally exploit the model's structural properties to analytically derive monotone decision policies, which complement the numerical results. Our model could help health care professionals in improving their counseling of mutation carriers since patients can provide individualized QOL preferences for surgery.

Our primary results recommend that BRCA1 carriers undergo BM starting at age 30 up until age 60, after which BM is no longer recommended, and undergo BSO from age 40 onwards to maximize cumulative QALYs in the baseline scenario. For BRCA2 carriers, the window for recommended BM is reduced to ages 40 to 46, after which it is no longer advised, and the optimal age to undergo BSO is delayed to age 49 and later, as the overall cancer risk is lower compared to BRCA1.

The remainder of the paper is organized as follows: Section 3.2 provides an overview of the existing literature on cancer risk for BRCA1/2 carriers, risk-reduction methods and the use of MDP models as decision aids in health care. The structure of our MDP model is described in Section 3.3, and methods for solving it are in Section 3.4. Numerical results of the optimization approach are presented in Section 3.5 and we conclude with discussion in Section 3.6.

3.2 Literature Review

3.2.1 BRCA and cancer

The link between early onset of breast cancer and genetic mutations was first discovered in 1990 for BRCA1 (Hall et al., 1990) and 1994 for BRCA2 (Wooster et al., 1994). In their meta-analysis for female mutation carriers aged 20 to 70 years old, Chen & Parmigiani (2007) report a mean cumulative breast cancer risk of 57% (BRCA1) and 40% (BRCA2), and ovarian cancer risk of 40% (BRCA1) and 18% (BRCA2). Clinical studies provide evidence that these risks can be reduced through performing a BM, BSO, or both. Breast cancer risk is reduced by a BM as there is minimal breast tissue remaining that that could develop cancer. McDonnell et al. (2001) report a 90% breast cancer risk reduction with a prophylactic BM, and Rebbeck et al. (2004) estimate a 94% reduction. A BSO reduces the risk of ovarian cancer by an estimated 80% through removing the organs themselves (Rebbeck et al., 2009). This procedure has the added benefit of lowering breast cancer risk by 30% to 59% depending on the age at surgery, due to reduced estrogen production, which can fuel breast cancer growth Eisen et al. (2005). Domchek et al. (2006) find a positive effect of prophylactic BSOs on overall survival as well as breast and ovarian cancer-specific mortality of mutation carriers. Their findings are in line with other clinical studies (Domchek et al., 2010; Kauff et al., 2008) and are confirmed by a meta-analysis (Rebbeck et al., 2009).

3.2.2 Quality-of-life

While the medical community mostly agrees on the risk-reducing benefits of a BM or BSO, there is less concordance regarding their respective impact on a patient's QOL, defined as a value between 1 (corresponding to perfect health) and 0 (death), of a patient's well-being given a specific condition or disease. For prophylactic BM, Grann et al. (1999b) and Grann et al. (2010) report a reduced long-term QOL based on an empirical evidence they find after administering questionnaires to women with an increased risk for breast cancer and a control group. At least a temporary lower QOL is suggested by Barton et al. (2005). They find that two-thirds of women undergoing a BM suffer from at least one, mostly reversible, complication. Several other studies find no long term effects of a BM. The two study populations examined by Tercyak et al.

(2007) consist of women with a BRCA1/2 mutation and unilateral breast cancer who chose a contralateral mastectomy (i.e., surgically removing the other non-cancerous breast) or breast conserving therapy (i.e., removing only the breast lump and nearby tissue) as a treatment option. Their results did not show a significantly lower QOL associated with a BM. Geiger et al. (2006) draw similar conclusions after comparing QOL between women who had been diagnosed with unilateral breast cancer and chose either a contralateral mastectomy or other treatment options.

The literature on QOL following BSOs also shows heterogeneous findings. Lower long-term health utilities are reported by Grann et al. (1999b, 2010). A review by Shuster et al. (2008) suggests that a BSO, especially in pre-menopausal women, increases the risk of negative side effects including osteoporosis, cardiovascular disease, and decline in sexual activity. Shuster et al. (2008) suggest that QOL may depend on the surgery's timing relative to the onset of natural menopause. Their review provied indications of a lower QOL after a BSO among pre-menopausal women, as the surgery eliminates fertility and immediately leads to surgically-induced menopause. Fry et al. (2001) did not find significant differences in the QOL of women who had undergone a prophylactic BSO and women who chose a screening program instead, although women report lower health utilities on some sub-scales after a BSO. Similar conclusions are drawn by Madalinska et al. (2005); Michelsen et al. (2009); Robson et al. (2003). For an extensive review of the literature in this area we refer to Harmsen et al. (2015).

3.2.3 Decision analysis models

Decision analysis using Markov Chains and Decision Process models have been applied to several health care topics, including decision support for BRCA1/2 carriers. Mason et al. (2011) and Mason et al. (2014) numerically optimize medication decisions for diabetes patients through solving a MDP with DP. Decision before making a liver transplant are analysed by Alagoz et al. (2004, 2007), who identify optimal policies for the timing of living-donor transplantations and the acceptance or rejection of cadaveric organs both through numerical results and structural insights into their models. Sandıkçı et al. (2008) determine the lost life expectancy for a patient on the liver transplantation waiting list through concealing information about the composition of the waiting list. Shechter et al. (2008) determine the QALY-maximizing time to initiate HIV treatment while Khademi et al. (2015) solve an MDP with ap-

proximate dynamic programming to determine the value of different HIV treatment prioritization practices.

Several papers use MDPs to provide recommendations for breast cancer screening schedules or treatment choices based on screening results. Ayer et al. (2012) use partially observable MDP to determine personalized mammography schedules based on a women's screening history and personal risk. Chhatwal et al. (2010) and Alagoz et al. (2013) both provide numerical results and structural properties of an MDP model to improve decisions about follow-up actions after a mammography exam. An overview about the application of MDP in health care is given by Alagoz et al. (2010).

Several studies employ Markov chains which model health states of BRCA1/2 mutation carriers to examine different prevention strategies and outcome measures. Kurian et al. (2012) provide a decision analysis tool that computes changes in mortality rates for different ages and both mutations based on the timing of surgeries. Schrag et al. (1997) use a Markov model with data from clinical studies to compare the resulting life expectancy of nine different prevention strategies, each varying either in the performed procedures or their timing. One of their findings suggests that delaying a BSO for 10 years at the age of 30 has limited impact on the life expectancy. Grann et al. (2002) and Schrag (2000) consider Tamoxifen therapy as an additional option, while Armstrong et al. (2004) assess the effect of hormone replacement therapy on the outcome of a BSO. In addition to measuring survival benefits, quality-adjusted outcome measures like QALYs are used by Grann et al. (1998) and van Roosmalen et al. (2002). Anderson et al. (2006) analyze the cost-effectiveness of different prevention strategies.

While Markov chains can provide an indication how different static decision policies alter certain outcome measures, the computation of optimal decision policies requires the us of an MDP in combination with a solution method, e.g. DP. Abdollahian & Das (2015) develop an MDP to compute the optimal timing of prophylactic BM and BSO to achieve a cost- or QALY-optimal strategy for mutation carriers. They find that for a BRCA1 carrier, the QALY-optimal strategy would be a BSO at age 30 and a BM at age 50. Although their model optimizes a carrier's accumulated QALYs, it does not track the development of the a carrier's health state after a cancer diagnosis and distinguishes between tumor subtypes based only on the estrogen receptor status, which lead to misestimates of the cancer-specific mortality rate.

The existing literature offers limited guidance to BRCA mutation carriers and physicians regarding the optimal timing of prophylactic surgeries. Studies either lack the ability to choose an optimal strategy from all relevant decision alternatives, or require assumptions that ignore important characteristics of breast and ovarian cancer progression. We contribute to the literature in three ways. First, we develop the most comprehensive decision analytic model for BRCA1/2 mutation carriers that is also computationally tractable and capable of finding the optimal surgery sequence. Second, we exploit the structural properties of a simplified model version and analytically derive monotone decision policies, which could help foster a better understanding of the model's results in certain circumstances. Third, we provide insights into the influence of a carrier's personal preferences on her QALY-maximizing sequence of prophylactic surgeries.

3.3 Model Formulation

We present a novel Markov Decision Process (MDP) model of the prophylactic surgery decision faced by BRCA1/2 mutation carriers: at what age(s) to undergo a BM and BSO to reduce her breast and ovarian cancer risk to maximize quality-adjusted life expectancy. For the full model, we use an finite horizon with yearly decision epochs.

We assume that carriers are initially healthy with no prior history of breast or ovarian cancer, have not previously undergone a BM or BSO surgery, and are first eligible for surgery at 20 years old. Model states correspond to various health states, including diagnosis of breast and ovarian cancer and associated tumor subtypes and stages. Undergoing a BM or BSO reduces the risk of breast or breast and ovarian cancer, respectively. The carrier's mortality rate and QOL depends on her age and health state. Should a cancer diagnosis occur, a cancer-specific mortality rate that depends on the tumor site (breast or ovaries), stage, and subtype (for breast cancer only) is added to the baseline mortality rate. Model parameters differ between the two mutation types wherever evidence from clinical studies could be found.

3.3.1 State space

The state space $S = \{age,\ surg,\ bc,\ bsu,\ bcts,\ oc,\ octs,\ de\}$ consists of eight variables as summarized in Table 3.1. Each state variable takes on one value from a finite set described below:

age: indicates the age of the carrier in 1-year increments;
$age \in \{20,\ 21,\ \ldots,\ AGE\}$. The maximum age AGE is set to 85 years as no studies differentiate cancer risk for carriers older than this age.

surg: indicates which surgeries a carrier has previously undergone;
$surg \in \{None,\ BSO < 40,\ BSO = 40,\ BSO = 41,\ \ldots,\ BSO > 49,\ BM,\ BM\&BSO\}$. The age at BSO is saved for surgeries occurring between ages 40 and 50, as earlier removal of the ovaries reduces lifetime estrogen exposure, resulting in a lower breast cancer risk (Eisen et al., 2005).

bc: indicates whether a carrier has been previously diagnosed with breast cancer;
$bc \in \{None,\ In\ treatment,\ \leq 5\ years\ ago,\ > 5\ years\ ago\}$. The carrier can have no history of breast cancer, currently be in treatment, or be in a post-treatment stage. The mortality rate for certain breast cancer sub-types (e.g., early stage triple-negative) sharply drops after five years and distant recurrences beyond five years post-cancer treatment are rare (Lee et al., 2011). The variable therefore indicates if at least five years have passed since the completion of breast cancer treatment.

bcts: indicates the breast cancer tumor stage at the time of diagnosis;
$bcts \in \{None,\ I,\ II,\ III,\ IV\}$. As with most clinical studies, we distinguish between tumor stages, where stage I tumors are <2cm and only in the breast; stage II tumors are 2-5cm or spread to lymph nodes; stage III tumors are >5cm or spread to the chest wall; and stage IV indicates cancer spread to other organs such as the liver, lungs, brain, or bones. A later stage at diagnosis indicates a less favorable prognosis and higher mortality rate (Edge & Compton, 2010).

bcsu: indicates the tumor subtype when breast cancer has been diagnosed;
$bcsu \in \{None,\ LuA,\ LuB,\ TN,\ HER2\}$. We use a common breast cancer classification scheme consisting of four sub-types: Luminal A (LuA), Luminal B (LuB), Triple-Negative (TN), and HER2, which categorize tumors based

on common combinations of estrogen-, progesterone-, and HER2-receptor statuses. (Sørlie et al., 2003). TN and HER2 sub-types tend to be faster growing tumors, leading to later stage at diagnosis, higher rates of metastatic disease and therefore higher mortality rates (Brown et al., 2008).

oc: indicates whether a carrier has been previously diagnosed with ovarian cancer; $oc \in \{None,\ In\ treatment,\ Post-treatment\}$. The carrier can have no history of ovarian cancer, currently be in treatment, or be post-treatment.

octs: indicates the ovarian cancer tumor stage at the time of diagnosis; $octs \in \{None,\ I,\ II,\ III,\ IV\}$. Stage I corresponds to tumors contained within the ovaries or fallopian tubes; stage II tumors can have spread to the uterus or other pelvic organs; stage III tumors have spread to the abdominal lining or lymph nodes; stage IV indicates cancer spread to organs including the spleen, liver, lungs, and others.

de: indicates if a carrier is alive, died from metastatic cancer or from other causes; $de \in \{Alive,\ Metastatic\ cancer\ death,\ Other\ death\}$.

3.3.2 Action space

The model's action space consists of four actions $A = \{W,\ BM,\ BSO,\ BM\&BSO\}$. If state *surg* indicates that no surgeries have been performed yet, a carrier can choose an action $a_t(s)$ from the following options: Wait, undergo a BM, undergo BSO, or undergo both procedures. If *su* indicates that an organ has already been removed (either prophylactically or following a cancer diagnosis), the respective surgery is excluded from the action space. Although women diagnosed with cancer ultimately decide whether to undergo a BM or BSO as part of treatment, we assume that treatment of breast cancer includes a BM; similarly treatment of ovarian cancer includes a BSO. These are reasonable assumptions for BRCA mutation carriers, as the risk of developing a new cancer in the future is extremely high for this population (Trainer et al., 2010).

Table 3.1: State space variables and corresponding set of possible values

State variable	Variable value set						
Age (*age*)	20	21	22	23	...	85	
Prior surgery (*surg*)	None	BSO < 40	BSO = 40	...	BSO > 49	BM	BSO & BM
Breast cancer (*bc*)	None	In treatment	≤ 5 yrs ago	> 5 yrs ago			
Tumor stage (*bcts*)	None	I	II	III	IV		
Tumor subtype (*bcsu*)	None	Luminal A	Luminal B	Triple Negative	HER2		
Ovarian cancer (*oc*)	None	In treatment	Post treatment				
Tumor stage (*octs*)	None	I	II	III	IV		
Death (*de*)	Alive	Metastatic cancer death	Other death				

3.3.3 State rewards and transition probabilities

A short-term reward, $r_t(s_t, a) \in [0, 1]$, is assigned to each state-action pair and deterministically determined by the carrier's health state at time t. If the action includes a surgery, this has either a temporary or lifelong impact on the state reward. If cancer is diagnosed, the negative effect on the reward is dependent on the cancer stage as indicated in Table 3.5. In each state, the minimum of the QOL factors corresponding to each state variable is chosen as the period reward. Therefore, cumulative QALYs are calculated by summing over all health states that are visited before death.

Let P_t denote the matrix of all transition probabilities at time t. The probability that a carrier transitions from state s_t to s'_{t+1} when choosing action a is denoted by $p_t(s'_{t+1} \mid s_t, a)$.

Let $v^*(s_t)$ denote the optimal expected future reward when a carrier is in state s_t. The value of future periods is discounted with factor γ, $\gamma \in (0, 1)$. We can find $v^*(s_t)$ by solving the following Bellman equation (Puterman, 2014):

$$v(s_t) = \max \left\{ r_t(s_t, a) + \gamma \sum_{s'_t \in S} p_t(s'_{t+1} \mid s_t, a) v(s'_{t+1}),\ \forall s_t \in S, a \in A(s_t) \right\}. \tag{3.1}$$

3.4 Solution Approach

As the full model described in Section 3.3 has eight state variables which consists of at least three possible state values and the action set contains up to four actions, developing an analytical solution is infeasible due to the high complexity of the problem. We therefore limit the problem's scope to only one cancer type and one possible surgery and analytically derive an optimal control limit policy in Section 3.4.1. The size of the model's state space also imposes the 'curse of dimensionality', making the computation of an optimal solution through DP intractable. In Section 3.4.2 we therefore present an approach based on LP which is able to obtain an optimal numerical solution.

3.4.1 Structural properties

In this section, we show that a threshold policy exists for patients who have already been diagnosed with breast or ovarian cancer (and thus undergone a BM or BSO, respectively), and then must decide if and when to undergo the other surgery (BSO or BM, respectively). This analysis complements the numerical solution presented in Section 3.5 through showing that in these specific circumstances, surgery is always optimal when the remaining cancer risk is sufficiently high. We first consider a breast cancer survivor needs to decide if she wants to perform a BSO after having already performed a BM. We present the reverse solution for an ovarian cancer survivor thereafter. The approach presented in this section has been extended from an approach in Chhatwal et al. (2010).

Following a breast cancer diagnosis and treatment, the action set A is subsequently reduced to $\hat{A} = \{W,\ BSO\}$, with a BSO as the only remaining surgery option. The full model state space is S is reduced from an eight-dimensional to a one-dimensional vector $\hat{S} = \{x^{20},\ x^{21},\ \ldots,\ x^{AGE},\ x^{surg},\ x^{oc},\ x^{death}\}$, where $x^{20},\ x^{21},\ \ldots,\ x^{AGE}$ represents the carrier's age-induced breast cancer risk; these states assume no surgery or ovarian cancer diagnosis. We merge the state variables oc and $octs$ into a single state x^{oc}, which indicates the diagnosis of ovarian cancer. The state x^{surg} indicates that a BSO surgery has been performed. Finally, x^{death} indicates the death of the carrier through metastatic cancer or other causes.

The state x^{death} is defined as absorbing state with $r_t(x^{death} \mid \cdot) = 0$. After having undergone a BSO, the action set is reduced to $\hat{A}(x^{surg}) = \{W\}$ and $r_t(x^{surg}, W)$ is reduced by the impact of the surgery on the QOL. After a breast cancer diagnosis, $r_t(x^{oc}, \cdot)$ is additionally reduced by the impact of the cancer itself. We furthermore assume the following boundary conditions when reaching T, the last decision epoch which is defined as $T = 64$ in our case: In T, the immediate reward of performing a BSO equals the reward of waiting in T which reflects the conditional discounted QALYs of a healthy mutation carrier at T. When already having performed the BSO, the terminal reward accounts for the reduced mortality through the cancer risk reduction as well as the QOL impact of the BSO. The terminal reward when being alive after a cancer diagnosis is defined alike as the conditional discounted QALYs of

a mutation carrier diagnosed with cancer who is alive at age T:

$$\begin{aligned} r_T(s_T, BSO) &= r_T(s_T, W) \\ &= v_T(s_T), \ \forall \ s_T \in \{x^{20}, \ x^{21}, \ \ldots, \ x^{AGE}\}, \\ r_T(x^{surg}, W) &= v_T(x^{surg}), \\ r_T(x^{oc}, W) &= v_T(x^{oc}). \end{aligned} \tag{3.2}$$

All proofs are given in the Appendix, Section B.1.2. We make the following assumptions:

Assumption 3.1. *The probability of an ovarian cancer diagnosis $p_t(x^{oc} \mid s_t, W)$ without undergoing BSO is nondecreasing in $s_t, s_t \in \{x^{20}, \ x^{21}, \ \ldots, \ x^{AGE}\}$.*

Several medical studies support this assumption that the annual risk of developing ovarian cancer among BRCA carriers increases or remains constant as women age (Chen & Parmigiani, 2007).

Assumption 3.2. *The immediate reward $r_t(s_t, W)$ is nonincreasing in $s_t \ \forall \ s_t \in \hat{S}$ and in t.*

This implies that a carrier always moves to an equal or lower QOL state as she ages, $s_t \in \{x^{20}, \ x^{21}, \ \ldots, \ x^{AGE}\}$; undergoes a BSO, $s_t \in \{x^{surg}\}$; is diagnosed with ovarian cancer, $s_t \in \{x^{oc}\}$; or dies, $s_t \in \{x^{death}\}$.

Assumption 3.3. *The function $v_t(s_t, BSO)$ is nonincreasing in t and in $s_t \forall s_t \in \{x^{20}, \ x^{21}, \ \ldots, \ x^{AGE}\}$.*

This implies that a carrier's post-BSO QALYs are not increasing with her age. A possible decrease in QALY impact of a BSO after the onset of natural menopause will likely be offset by the decrease in remaining QALYs while aging.

Assumption 3.4. *The function $v_t(x^{oc}, \cdot)$ is nonincreasing in t.*

This implies that a carrier's expected remaining QALYs after an ovarian cancer diagnosis are not increasing with her age.

Assumption 3.5. *The function $v_t(x^{surg}, \cdot)$ is nonincreasing in t.*

This implies that a carrier's expected remaining QALYs after a BSO are not increasing with her age.

Assumption 3.6. $v_t(s_t, BSO) \geq v_t(oc, BSO) \ \forall \ s_t \in \{20, \ 21, \ \ldots, \ AGE\}$.

This means that expected remaining QALYs following a BSO are at least as great as expected remaining QALYs after ovarian cancer.

Assumption 3.7. *The transition probabilities for action W satisfy the following:*

$$p_t(x^{oc} \mid i, W) \leq p_{t+1}(x^{oc} \mid i, W), \tag{3.3}$$

$$p_t(x^{death} \mid i, W) \leq p_{t+1}(x^{death} \mid i, W), \tag{3.4}$$

for $i, j \in \hat{S} \backslash \{x^{surg}\}$.

This assumption means that the cancer risk is nondecreasing in t; as a carrier ages, she is more likely to transition to a worse health state (i.e., ovarian cancer or death) if she chooses to not undergo a BSO.

Definition 3.1. *Following Barlow & Proschan (1965), a Markov chain is assumed to be of increasing failure rate (IFR), if its rows are in increasing stochastic order, as:*

$$q(i) = \sum_{j=k}^{x^{death}} p(j \mid i) \tag{3.5}$$

is nondecreasing in $i \ \forall \ k \in \{x^{20}, \ x^{21}, \ \ldots, \ x^{death}\}$.

The above definition means that for the underlying Markov Chain of the MDP, as a carrier progresses into states with a higher cancer risk, her risk of progressing to states with an even higher risk of death also increases.

To show the existence of an optimal control limit for performing a BSO, we first show the monotonicity of $v_t(s_t)$ in s_t and t. We refer to the Appendix, Section B.1.1 for the respective Propositions and Lemmas.

Theorem 3.1. *If the transition probability matrix P_t for action W is IFR and satisfies the following two conditions:*

$$\frac{v_t(s_t, BSO) - v_t(s_t + 1, BSO)}{\gamma v_{t+1}(s_t + 1, BSO)} \leq p_t(x^{death} \mid s_t + 1, W) - p_t(x^{death} \mid s_t, W), \tag{3.6}$$

$$\sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t+1, W) \geq \sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t, W), \tag{3.7}$$

for all $s_t \in \{x^{20}, x^{21}, \ldots, x^{AGE}\}$ *and* $t = 1, 2, \ldots, T-1$*, then there exists an optimal control threshold state* $\overline{s_t} \in \{x^{20}, x^{21}, \ldots, x^{AGE}\}$ *for* $t = 1, 2, \ldots, T-1$ *such that:*

$$a^*(s_t) = \begin{cases} W & \text{if } s_t < \overline{s_t}, \\ BSO & \text{if } s_t \geq \overline{s_t}. \end{cases} \tag{3.8}$$

Inequality (3.6) denotes that as a carrier's age increases, the *decrease* in QOL due to a BSO is less than the *increase* in mortality risk by waiting one additional period. Inequality (3.7) requires that the probability of moving to a higher risk state or death increases with age.

An equivalent structural property can be derived for a carrier's decision to undergo BM after an ovarian cancer diagnosis and treatment, resulting in the following Theorem:

Theorem 3.2. *If the transition probability matrix* P_t *for action* W *is IFR for and satisfies the following condition:*

$$\frac{v_t(s_t, BM) - v_t(s_t+1, BM)}{\gamma v_{t+1}(s_t+1, BM)} \leq p_t(x^{death} \mid s_t+1, W) - p_t(x^{death} \mid s_t, W), \tag{3.9}$$

$$\sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t+1, W) \geq \sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t, W), \tag{3.10}$$

for all $s_t \in \{x^{20}, x^{21}, \ldots, x^{AGE}\}$ *and* $t = 1, 2, \ldots, T-1$*, then there exists an optimal control threshold state* $\overline{s_t} \in \{x^{20}, x^{21}, \ldots, x^{AGE}\}$ *for* $t = 1, 2, \ldots, T-1$ *such that:*

$$a^*(s_t) = \begin{cases} W & \text{if } s_t < \overline{s_t}, \\ BM & \text{if } s_t \geq \overline{s_t}. \end{cases} \tag{3.11}$$

We provide numerical results for Theorems 3.1 and 3.2 in Section 3.5.

3.4.2 Solution via linear programming

Most of the MDP's states have a limited number of feasible successor states, resulting in a sparse transition matrix. Solving the problem in Equation (3.1) through value iteration, the resulting DP would calculate the value of all states, of which most are never visited. We exploit the sparsity of the transition matrix to solve the MDP using linear programming (LP), an approach that would result in unfavourable computational performance for non-sparse transition matrices (Puterman, 2014). We obtain an optimal solution following an approach presented in White & White (1989) and de Farias & van Roy (2003).

We switch from a finite- to infinite-horizon MDP, as the LP approach is only capable of solving discounted infinite-horizon MDPs. In the finite version, the probability that a carrier significantly older than T will arrive in absorbing state *Other death* is high: $p(Other\ death(s_{t+1}) \mid s_t) \approx 1\ \forall\ t \gg AGE$. We therefore do not expect differences between the finite and infinite version, as the reward for a state indicating death is zero: $r_t(s_t, a) = 0\ \forall\ a \in \hat{A}$ and $\forall\ s_t :\ de(s_t) = Other\ death$.

An optimal policy $\pi*$ which satisfies (3.1) is given by:

$$\pi^*(s) = \underset{a \in A}{\operatorname{argmax}} \left\{ r(s,a) + \gamma \sum_{s' \in S} p(s' \mid s, a) v(s') \right\}. \tag{3.12}$$

Powell (2011) shows that at optimality v is smallest value that satisfies inequality:

$$v \geq \max_{a \in A} \left\{ r(s,a) + \gamma \sum_{s' \in S} p(s' \mid s, a) v(s') \right\}. \tag{3.13}$$

To find $v_s^* \forall s \in S$, we therefore need to minimize $v(s) \forall s \in S$ while imposing Equation (3.13) as constraint for every state-action pair. We can now formulate the linear program:

$$\min_{v} \quad \sum_{s \in S} v(s) \tag{3.14}$$

subject to

$$v(s) \geq r(s,a) + \gamma \sum_{s' \in S} p(s' \mid s, a) v(s'), \quad \forall s \in S, a \in A \tag{3.15}$$

After having calculated the optimal state values, we can calculate a corresponding

optimal policy through Equation 3.12.

Due to the sparsity of the transition matrix, the number of transitions which are not zero is significantly smaller than a full transition matrix, making the problem tractable. The optimization of the MDP from Section 3.3 results in an LP with 7,775,625 constraints and 4,387,500 variables. The underlying algorithms are coded in Matlab, R2015b, and the LP is solved by Gurobi 6.5.1. To ensure the applicability of the optimization results to the research question, we ensure the validity of the model through using an approach proposed by Gass (1983) and Sargent (2013). Detailed results are presented in the Appendix, Section B.3.2.

3.5 Numerical Study

Although we analytically develop insights into the structure of a simplified version of the model where a carrier is only considering one surgery type (Section 3.4.1), we must make some simplifying assumptions in order to determine an optimal threshold policy (Theorems 3.1 and 3.2). In particular, we ignore the effect of a prophylactic BSO on breast cancer risk, caused through a lower estrogen production. The optimal policy does not capture the complexity of the full model with multiple surgery options (Section 3.3). We therefore provide numerical results for the full model in the remainder of this Section.

We conduct various numerical analyses to provide insights into optimal surgery timing and sequence under differing assumptions. We vary the duration and magnitude of a surgery's impact on a carrier's QOL. To best reflect potential decisions faced by different BRCA mutation carriers, we also consider constraining the minimum eligible age for a BSO (i.e., to mimic a carrier's preference to delay the surgery until after childbearing), and we examine optimal strategies after either breast or ovarian cancer has been previously diagnosed. To ensure the validity of our model, we follow a validation process outlined in Section B.3.

3.5.1 Data sources

The model's transition probabilities and rewards are based on values obtained from published clinical studies. We use BRCA1- and BRCA2-specific values when possible. If multiple data sources are available, we select those with greater sample sizes and the most recent studies to reflect the latest advances in cancer therapy. For the baseline risk for breast- and ovarian cancer we use the results of a meta-analysis by Chen & Parmigiani (2007).

Key probabilities include age-dependent breast and ovarian cancer risk; distribution of each cancer by stage and sub-type (breast cancer only); and cancer risk-reduction resulting from a prophylactic BM or BSO (Table 3.3 shows the distribution of subtypes and stages for cancer diagnoses as well as the risk reduction through each prophylactic surgery). Cancer-specific mortality rates used in the model are given in Table 3.4, and age-dependent mortality rates from other causes are obtained from Centers for Disease Control and Prevention & National Center for Health Statistics (2014). To reflect current practices in medical treatemnt, no studies older than 2010 are considered for mortality rates. Quality-of-life values for each health state used in the model are presented in Table 3.5. Unless otherwise stated, the model's objective is to maximize a carrier's discounted lifetime QALYs.

3.5.2 QALY-maximizing surgery sequence

3.5.2.1 BRCA1

The QALY-maximizing prophylactic surgery sequence for a healthy (no personal cancer history) BRCA1 carrier results in a life expectancy of nearly 77 years, compared to 70 years with no prophylactic surgery, or 73 years if surgeries are delayed until after age 50. Although life expectancy 0.4 years longer if both surgeries are performed very early at the age of 30, cumulative QALYs are lower with early surgeries. The QALY-maximizing sequence balances this trade-off of survival benefit with early surgery against diminished QOL for the carrier following a BM and BSO.

Under the baseline parameter assumptions, the optimal surgery sequence suggests first undergoing a BM between ages 30 and 60. Before age 30, the very low risk of breast

cancer does not justify the loss in QOL through a BM; after age 60, the remaining breast cancer risk is not great enough to justify a BM, especially because a BSO can also reduce breast cancer risk. In other words, undergoing a BM incurs a "fixed cost" because we assume a 5-year hit to QOL, which must be amortized over a sufficient number of years in the form of reduced breast cancer risk to warrant the upfront cost.

The optimal policy then recommends a BSO for any BRCA1 carrier over age 40, because her risk for ovarian cancer is sufficiently high to justify the expected reduction in QOL. This surgery is recommended until age 85, as the 1-year impact on QOL in post-menopausal women over age 50 is very small compared to the risk reduction of both breast and ovarian cancer. Figure 3.1 depicts the optimal surgery sequence for every age ranging from 20 years to 85.

We also use our MDP model to simulate the incidence of breast and ovarian cancer under different prophylactic surgery policies (Figure 3.2). For BRCA1 carriers, the simulated risk of breast cancer by age 65 decreases from 49.1% (with no surgeries) to 5.8% with the QALY-maximizing policy, and ovarian cancer decreases from 30.3% (with no surgeries) to 9.4% with the QALY-maximizing policy. Although the risks of cancer (5.1% for breast, 8.1% for ovarian) are not quite as low as the more aggressive policy of undergoing both BM and BSO at age 30, the QALY-maximizing provides 94-98% of the cancer-reducing benefits of an aggressive policy, but importantly considers both the overall QOL and survival gains from such dramatic surgeries. Nevertheless, carriers who wish to minimize their cancer risk as much as possible may prefer to undergo a BM and BSO at an earlier age.

3.5.2.2 BRCA2

For BRCA2 carriers, the optimal sequence delays both surgeries due to the lower overall risk of breast and ovarian cancer. Undergoing a BM is only recommended for carriers aged 40 to 46 years. Afterwards, the reduction in the future risk of breast cancer (assuming a carrier has remained cancer-free until age 46) does not offset the loss in QOL imposed by a prophylactic BM. However, beginning near age 50, a BSO is recommended thereafter, as a carrier would otherwise enter natural menopause around this age.

Under the QALY-maximizing surgery sequence, life expectancy increases to 78 years

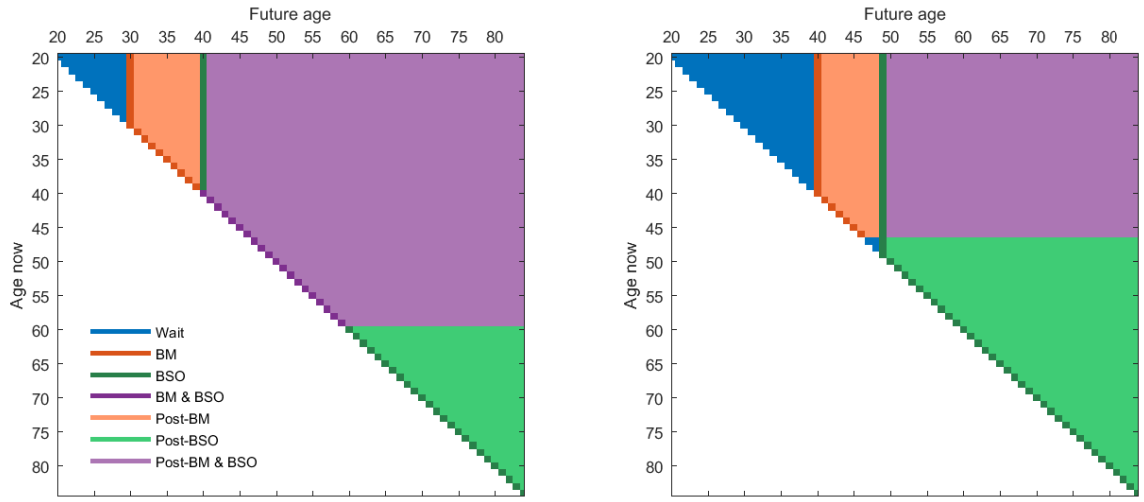

Figure 3.1: QALY-maximizing surgery sequence for BRCA1 (left) and BRCA2 (right) carriers who have not previously developed breast or ovarian cancer, nor undergone prophylactic surgery.

from 75 years with no surgeries, a smaller overall gain than for BRCA1 carriers because the overall risk of cancer is lower and thus the marginal benefits of undergoing prophylactic surgery are more modest (Table 3.2).

As with BRCA1, the QALY-maximizing policy results in a higher incidence of breast cancer compared with a strategy of undergoing BM and BSO at age 30, because the QALY-maximizing policy recommends a BM only for a limited group of carriers aged 40 to 46. The risk of ovarian cancer by age 65 is reduced from 11.3% to 4.6%, which is slightly higher than the 2.4% if women elect to have both surgeries at age 30.

3.5.2.3 Sensitivity analysis

While parameters governing the distribution of cancer stages and sub-types, surgery risk-reduction, and cancer-related mortality rates are broadly consistent across multiple clinical studies, there exists variability in estimates of QOL following prophylactic surgery in those at high-risk of developing cancer. These variations are largely attributable to differences in study populations (e.g., mutation carriers, cancer survivors, health care professionals) or methodologies at eliciting preferences.

As both surgeries impact physical and psychological well-being, we examine how dif-

Table 3.2: Average life expectancy, survival probability, and cumulative risk of breast and ovarian cancer under different decision policies.

Mutation type	Decision policy	Average life expectancy	Survival prob. by age 85	Cumulative cancer risk by age 65	
				Breast	Ovarian
BRCA1	Without surgeries	69.8	60.6%	49.1%	30.3%
	QALY-max.	76.8	81.6%	5.8%	9.4%
	BM at 30	73.1	71.0%	7.5%	30.0%
	BSO at 30	75.3	77.1%	25.2%	8.4%
	BM & BSO at 30	77.2	82.8%	5.1%	8.1%
BRCA2	Without surgeries	75.4	77.1%	40.7%	11.3%
	QALY-max.	77.8	83.9%	10.2%	4.6%
	BM at 30	77.3	82.8%	5.4%	11.6%
	BSO at 30	77.9	84.1%	19.3%	2.7%
	BM & BSO at 30	79.0	87.1%	3.5%	2.4%

BM = Bilateral mastectomy (surgical removal of both breasts).
BSO = Bilateral salpingo-oophorectomy (surgical removal of both ovaries and fallopian tubes).

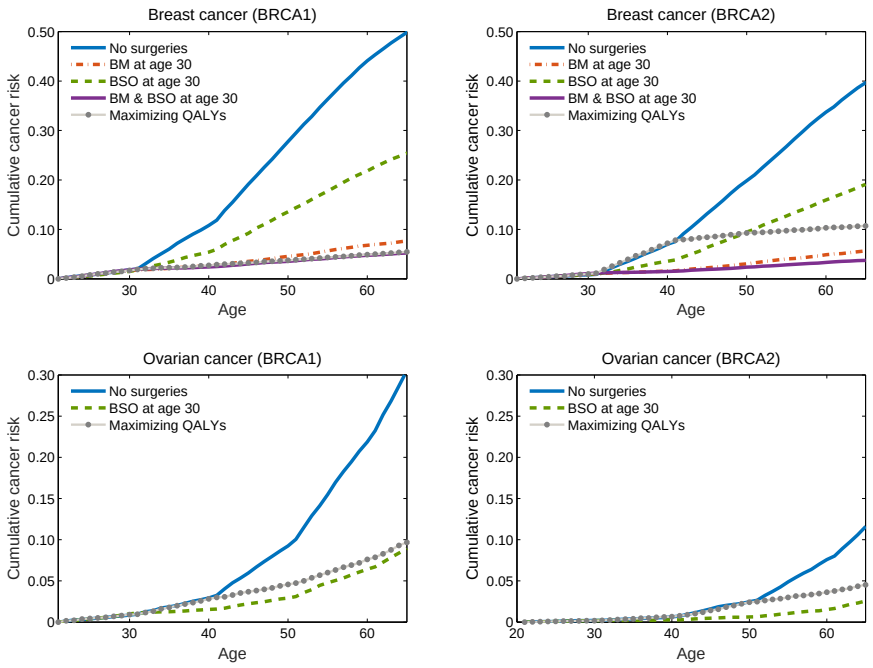

Figure 3.2: Cumulative cancer risk for different surgery sequences.

ferent assumptions regarding QOL, reflecting different patient preferences, affects the optimal surgery sequence (Figure 3.3). For BRCA1 carriers, lowering the impact of

a BM by 50% would expand the window to undergo the surgery to up until age 77. For a BSO, a 50% reduction in its impact on QOL flips the sequence of surgeries, recommending BSO starting at age 30 followed by BM from age 40 to 60, reflecting the lower impact of a BSO. When varying QOL impact of both surgeries, results show that different preferences about the impact of a BM have a higher effect on the decision if and when it should be performed, compared to varying preferences about a BSO impact (see Figure 3.3 and 3.4). This is likely due to the fact that a BSO simultaneously reduces the risk for both cancers, while a BM has a static impact on the QOL, which has to be compensated during the remaining lifespan. For ages 30 to 39 and a BM 1-year QOL impact of 0.5, our results furthermore indicate that it is QALY-maximizing to perform a BSO in order to avoid breast cancer and subsequently the high QOL impact of the BM performed during treatment.

For BRCA2 carriers, we find qualitatively similar results as for a 50% reduction of the impact of a BM, the youngest age at which a BM is QALY maximizing is lowered from 40 to 30. When reducing the impact of a BSO by 50%, the youngest age at which a BSO is recommended is lowered from 50 to 40, and the surgery sequence does not include a BM at any age as most of the breast cancer risk, also for carriers in their 40s, has been removed by the BSO.

3.5.3 Constrained surgery options

3.5.3.1 Prior cancer diagnosis

Some mutation carriers might consider a limited set of surgery options, either due to personal preferences or to a prior cancer diagnosis that required one of the surgeries to be performed as part of the treatment, as it is often the case for BRCA mutation carriers (Trainer et al., 2010). Following a diagnosis of ovarian cancer, a BSO as part of treatment also reduces subsequent breast cancer risk; therefore, also undergoing a BM is no longer QALY-maximizing at any age. In a similar vein, after a breast cancer diagnosis and bilateral mastectomy, the marginal benefits of undergoing BSO on breast cancer are diminished. Among young breast cancer survivors who have undergone a BM, the optimal timing of BSO increases from age 40 to 46 (BRCA1) and from age 43 to 48 (BRCA2). Beyond these ages, it is always QALY-maximizing to perform a BSO, numerically in line with Theorem 3.2, which states that performing a

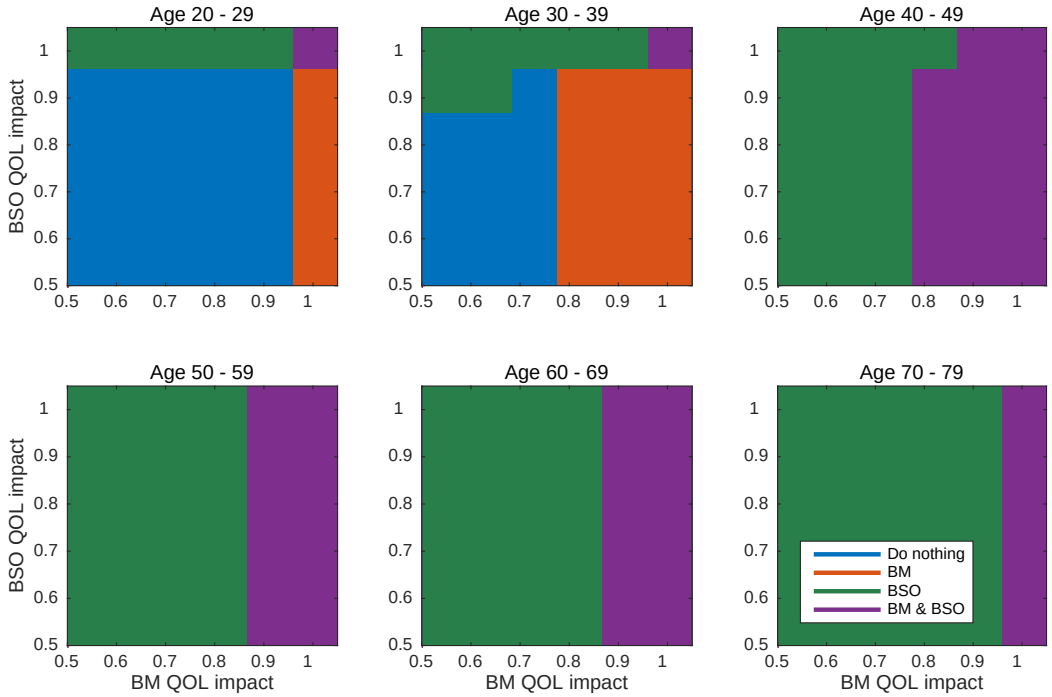

Figure 3.3: Optimal surgery sequence for BRCA1 carriers at age 20, 30, 40, 50, 60, 70 assuming varying health utilities for prophylactic as well as treatment-induced BM and BSO.

surgery is always optimal when the cancer risk exceeds the optimal control threshold $\overline{s_t}$.

To further illustrate this analysis, suppose a carrier is diagnosed with breast cancer at the age of 35 and undergoes a BM as part of treatment. If she elects to not undergo any additional prophylactic surgery, her life expectancy is approximately 59 years (BRCA1) or 61 years (BRCA2). However, if she decides to undergo a BSO at the QALY-maximizing recommended age, her life expectancy increases to 70 years (BRCA1) or 72 years (BRCA2), a substantial gain. Her ovarian cancer risk by age 65 is reduced from 30.1% to 11.2% (BRCA1) or from 11.5% to 4.1% (BRCA2) with this strategy.

3.5.3.2 Fertility considerations

Some women may prefer to delay surgery until a later age, especially a BSO, as the procedure eliminates future reproduction. They may wish to complete family planning or wait until the natural onset of menopause. As our model with baseline parameters

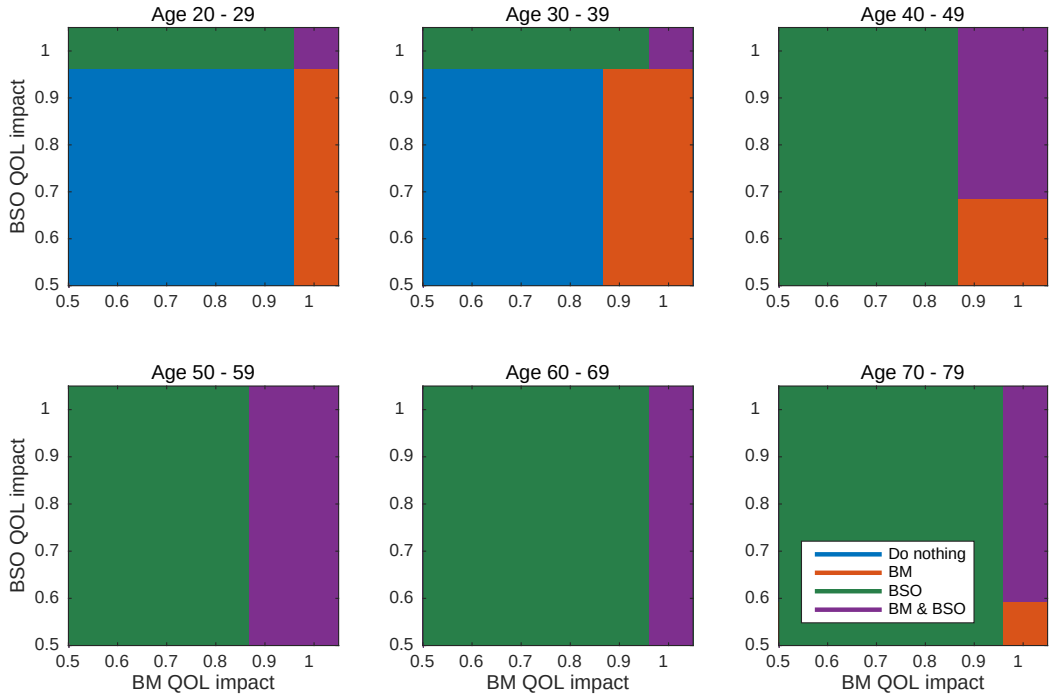

Figure 3.4: Optimal surgery sequence for BRCA2 carriers at age 20, 30, 40, 50, 60, 70 assuming varying health utilities for prophylactic as well as treatment-induced BM and BSO.

does not recommend performing a BSO before age 40 (BRCA1) or 50 (BRCA2), most women will not be substantially limited in their family planning through this recommendation. After having performed a BM at the age of 30, delaying a BSO until age 50 for BRCA1 carriers decreases the overall life expectancy by 1 year compared to the QALY-maximizing sequence, ovarian cancer risk by age 65 increases by 5.8 percentage points.

3.5.4 Improved breast cancer screening

Routine screening for ovarian cancer, unfortunately, has very limited accuracy (van Gorp et al., 2011), resulting in more frequent stage III or IV diagnoses for ovarian cancer than for breast cancer. In contrast, frequent mammography and Magnetic Resonance Imaging (MRI) may be effective alternatives to surgery for BRCA mutation carriers, to decrease breast cancer-specific mortality. Although enhanced screening will not *prevent* cancerous tumors from developing, it allows for earlier detection and treatment, improving the prognosis and survival. Our baseline parameter assumptions

reflect, to the best of our ability, the current screening practices among BRCA carriers, in particular, the stage at which patients typically present at diagnosis (Table 3.3). Investigating the impact of more or less frequent screening is important, as screening adherence is known to be low even for women at high-risk of breast cancer (Garcia et al., 2014).

We examine how changes in breast cancer screening efforts affect the optimal surgery sequence, under the optimistic assumption that breast cancer is always diagnosed at stage I. Compared to the original policy, a more conservative surgery sequence is recommended: BSO at age 39 (BRCA1) or age 49 (BRCA2), and no BM at any age (both BRCA1 and BRCA2). Of note, the probability of a breast cancer diagnosis increases by 27.7 percentage points compared to the baseline scenario, and life expectancy increases by one year as stage I tumors have minimal impact on cancer-specific mortality. Although this scenario would require essentially perfect screening accuracy and total adherence to the screening schedule, it indicates that increased screening efforts are a partial substitute for a prophylactic BM.

Table 3.3: Distributions of cancer sub-types, stages, and risk-reduction of prophylactic surgeries

Parameters [Unit]	BRCA1	BRCA2	Source
Breast cancer sub-types [%]			
Luminal A	04	10	Mavaddat et al. (2012)
Luminal B	22	72	
HER2	06	03	
Triple Negative	69	16	
Breast cancer stage at diagnosis [%]			
Luminal A/B			Brown et al. (2008)
I	49		
II	40		
III	07		
IV	03		
HER2			
I	29		
II	46		
III	18		
IV	07		
Triple Negative			
I	34		
II	50		
III	12		
IV	04		
Ovarian cancer stage at diagnosis [%]			
I	12	09	Bolton et al. (2012)
II	10	06	
III	64	73	
IV	13	12	
Breast cancer RR through BSO [OR[1]]			
≤ 40 years	0.41		Eisen et al. (2005)
40 – 50 years	interpolated		
≥ 50 years	0.70		
Ovarian cancer RR through BSO [HR[1]]	0.21		Rebbeck et al. (2009)
Breast cancer RR trough BM [HR[1]]			
BM only	0.09		Rebbeck et al. (2004)
BM & BSO	0.05		

BM = Bilateral mastectomy (surgical removal of both breasts).
BSO = Bilateral salpingo-oophorectomy (surgical removal of both ovaries and fallopian tubes).
HR = Hazard ratio (The ratio of the cancer rate with and without prophylactic surgery).
OR = Odds ratio (The number of cancer events after prophylactic surgery in a period divided by the number of events in the same period assuming no prophylactic surgery).
[1] See Appendix, Section B.2 for HR and OR conversion method used.

Table 3.4: Cancer-specific mortality rates

Tumor sub-type and stage	Annual mortality rate		Source
	BRCA1	BRCA2	
Breast cancer			
Luminal A, stage*			
I		0.004	
II		0.012	
III–IV		0.041	
Luminal B, stage			
I		0.008	
II		0.026	
III–IV		0.074	Parise &
HER2, stage			Caggiano
I		0.011	(2014)
II		0.037	
III–IV		0.099	
Triple Negative, stage			
I		0.015	
II		0.047	
III–IV		0.147	
Ovarian cancer			
I–II	0.06	0.04	Bolton et al.
III–IV	0.14	0.10	(2012)

* Breast cancer sub-types Luminal A and Luminal B, HER2- from Parise & Caggiano (2014) are combined into Luminal A, to match the sub-type definition used by Sørlie et al. (2003).

Table 3.5: Baseline quality-of-life (QOL) values of each health state

Health state	QOL factor	Impact duration	Source
Breast cancer			
Stage I-III	0.87	1 yr	Grann et al. (2010)
Stage IV (de novo)	0.59	Lifetime	Grann et al. (1999b)
Stage IV (recurrence)			
Luminal A	0.59	2.2 yr	Grann et al. (1999b), Lobbezoo et al. (2015)
Luminal B	0.59	1.6 yr	
HER2	0.59	1.3 yr	
Triple Negative	0.59	0.7 yr	
Ovarian cancer			
Stage I-III	0.84	1 yr	Grann et al. (2010)
Stage IV (de novo)	0.59	Lifetime	Grann et al. (1999b)
Stage IV (recurrence)	0.59	2.5 yr	
Surgery			
BM	0.88	5 yr	Grann et al. (2010)
BSO	0.95	Constant until age 50, 1 yr after age 50	

De novo = patient is found at stage IV cancer at diagnosis.
Recurrence = patient progressed from stage I, II, or III to stage IV.
BM = Bilateral mastectomy (surgical removal of both breasts).
BSO = Bilateral salpingo-oophorectomy (surgical removal of both ovaries and fallopian tubes).

3.6 Discussion

This paper presents a comprehensive approach to determine a recommended course of action for BRCA1/2 mutation carriers deciding if and when to undergo cancer risk-reducing prophylactic surgery. Our study develops a novel Markov Decision Process model of two simultaneous diseases—and multiple actions that can attenuate the incidence of both diseases—an advance not previously explored in the decision modeling literature. Our study combines an MDP model with a carefully appraised set of parameters based on observed clinical data. Although the size of the state space exceeds four million states, we are able to exploit the sparsity of the transition probability matrix and find an exact optimal solution using linear programming.

BRCA1/2 carriers are at substantially increased risk of breast and ovarian cancer compared to the general population. Our primary model findings indicate that to maximize a BRCA1 carrier's lifetime discounted quality-adjusted years, a BM is recommended between age 30 and 60, along with a BSO from age 40 onwards. Under

this sequence, our model computes an average life expectancy of 77 years for a 20-year old BRCA1 carrier, a 10% increase compared to performing no prophylactic surgeries. Results differ for BRCA2 carriers, with a later recommended age to undergo surgery breast and ovarian cancer risk is lower compared to BRCA1.

As preferences about the impact of prophylactic surgeries on a carrier's QOL most likely vary between carriers, we provide a sensitivity analysis demonstrating how a QALY-maximizing sequence would change when each individual surgery triggers a higher or lower impact on QOL. Results show that surgery schedules are more sensitive to changes in QOL impact of a BM than a BSO. We therefore conclude that under our assumptions, the surgery timing of Angelina Jolie (BRCA1, BM at age of 38, BSO in the early 40s), who's public appearance on that matter increased awareness about BRCA mutations, is in line with our model's results.

3.6.1 Limitations

As with any stylized mathematical model of an underlying disease process, our MDP model has several limitations. We select an appropriate state space to reflect key differences in transition probabilities, mortality, and QOL, and to ensure tractability in obtaining an optimal solution. Of course, this necessarily simplifies the complex tumor development and progression process, and nuances in different treatment outcomes.

With only one in 400 women carrying a BRCA mutation, many of whom are unaware of their status, most clinical studies involving BRCA carriers typically include small samples, often limiting the granularity and explanatory power of the data. We therefore are not able to obtain BRCA1- or BRCA2-specific estimates of the risk-reducing effects of surgeries. Although the baseline cancer risk varies between mutation types, the underlying mechanisms generating the risk reductions are likely equal for both mutation types, enabling us to use the same parameter values. Cancer-specific mortality rates are typically aggregated at 5-year or 10-year intervals, leading us to assume a constant mortality rate over time.

Structural limitations include a static screening policy and the omission of a local recurrence (to the breast or ovary) and other treatment options. Although the quality of breast cancer screening and adherence to the recommended schedule varies from woman to woman, we assume a constant screening rate, limiting the accuracy of the

cancer stage distribution at diagnosis. We try to overcome this issue through varying the distribution during sensitivity analysis.

Chemoprevention of breast cancer through medication, like Tamoxifen or oral contraceptives, is not included in the action space, resulting in a potentially incomplete representation of a carrier's choice set. We have excluded this therapy option as a strong risk reducing effect for all age groups and mutation types is uncertain (Duffy & Nixon, 2002; King et al., 2001). The risk of breast cancer recurrence or a contralateral breast cancer is significant for mutation carriers (Graeser et al., 2009; Nilsson et al., 2014). As the state space structure does not account for recurrences nor contralateral breast cancer and these states would alter a carrier's mortality, our model is limited by the accuracy of the breast cancer mortality rate. We overcome this limitation in first assuming bilateral mastectomies in the case of a breast cancer diagnosis avoiding contralateral cancer, and second capturing the implications of distant recurrences through the overall breast cancer mortality rate.

Another limiting factor is the magnitude and duration of the impact of prophylactic surgeries on women's QOL. Empirical studies present contradicting evidence for the impact of a BM and BSO (Harmsen et al., 2015). Most important, women exhibit varying preferences about surgery feasibility and timing, reflecting individual differences in family planning, self-perception, and perceived risk of cancer. We do not explicitly model the decision to undergo other procedures, including breast reconstruction following a bilateral mastectomy or oocyte retrieval prior to a bilateral salpingo-oophorectomy. To account for variability in these parameter values, we conduct an extensive sensitivity analysis varying both the impact of a BM and BSO.

3.6.2 Conclusions and future research

Through examining the structural properties of a simplified MDP model, we find that under the reasonable assumption that the cancer risk increases with age, an optimal-control limit exists, after which surgery is always QALY-maximizing. However, this analysis is limited to women who have already undergone either a BM or BSO, and it further assumes a minimal effect of a BSO on breast cancer risk.

We numerically confirm this optimal threshold for a BRCA carrier, who after having already been diagnosed with breast cancer is recommended to perform a BSO no

later than age 46 (BRCA1) or age 48 (BRCA2). Improved breast cancer screening adherence also impacts the recommended surgery schedule—under the optimistic assumption that *all* breast cancer tumors are diagnosed in stage I, our model does not recommend performing a BM at any age as the impact on QOL is not offset by a sufficient risk reduction.

Our model could be extended in several ways. By combining it with an underlying cancer model the consequences on QOL and the cancer specific mortality rate could become more precise. By modeling the QOL impact as a stochastic variable, the resulting decision policy would be more robust to uncertain outcomes of the surgeries (e.g., caused by complications). With more women seeking genetic testing for cancer-causing mutations, such as BRCA1/2, decision support models such as the one we have developed can help medical patients make better decisions under uncertainty, to improve both the length and quality of life. It also provides the basis for a possible cost effectiveness analysis which could be helpful to health care providers in assessing the optimal surgery sequence for cost effectiveness against the treatment of a possible breast or ovarian cancer.

Chapter 4

Spatial allocation of resources for epidemic control: A greedy heuristic versus dynamic programming approach

with Elisa F. Long and Stefan Spinler

The content presented in this chapter is based on Nohdurft et al. (2016a). Partial results from this chapter contributed to the conference presentations Nohdurft et al. (2015a) and Nohdurft et al. (2015b).

Allocating intervention resources effectively is vital for controlling geographically diverse and rapidly evolving infectious disease epidemics, such as the 2014 West Africa Ebola outbreak. We develop a two-stage model for optimizing when and where to assign Ebola treatment unit (ETU) beds during an outbreak's early phases. The first stage includes an epidemic model that forecasts occurrence of new cases at the regional level. The second stage includes two approaches to efficiently allocate intervention resources across affected regions: a greedy heuristic which prioritizes bed allocation based on a region's reproductive number which delivers good results when cased data are limited; and an approximate dynamic programming algorithm which utilizes the developed epidemic model to consider the impact of allocated resources and can be used when more nuanced data are available.

4.1 Introduction

Emerging infectious disease outbreaks, such as the 2014 Ebola epidemic in West Africa, or the ongoing Zika virus epidemic in Latin America, demonstrate the need for rapid and coordinated response efforts to mitigate the epidemics' spread. Before deploying critical medical resources—such as Ebola treatment units (ETUs), medications, and health care personnel—an essential step to ensure maximal benefit is assessing when and where such resources are most needed and anticipated to be needed. As an epidemic evolves temporally and geospatially, it becomes increasingly difficult to efficiently contain transmission. Achieving the most effective response requires a dynamic policy that integrates real-time updates in epidemiologic data, and adjusts intervention efforts to re-optimize resource allocation.

Mathematical models of Ebola transmission typically aggregate on a national scale (Gomes et al., 2014), despite evidence that Ebola—unlike other widespread pandemics, such as influenza—has remained concentrated in specific regions. The 2014 Ebola outbreak heterogeneously affected 51 out of 63 geographic regions in Liberia, Guinea, and Sierra Leone. Case numbers in affected regions ranged from fewer than 10 to more than 2,000, with one-third of regions accounting for 80% of all cases (Humanitarian Data Exchange, 2015a), suggesting that contagion could best be reduced with a targeted response. Physical proximity to regions with sizable infected populations undoubtedly affects both the spread of disease (Merler et al., 2015) and the ability to coordinate response efforts.

In mid-2014, some often cited models projected that nearly 600,000 new Ebola cases would occur by the outbreak's end, and as many as 1.4 million cases if underreporting were corrected (Butler, 2014b; Meltzer et al., 2014). Yet by the end of 2015, approximately 28,600 Ebola cases had been reported in West Africa (Humanitarian Data Exchange, 2015a). Early projections greatly overestimated new cases, in part because the underlying models assumed minimal or no behavioral changes in response to the epidemic's growth (Heesterbeek et al., 2015). Such models are typically calibrated to early transmission data and, when applied to later stages, their static parameters can lead to significant mis-estimates of disease incidence or prevalence. Once the risk of contracting Ebola was widely recognized in mid- to late-2014, transportation to and from the most affected regions was greatly reduced while social contact decreased locally. Our model accounts for this time-dependent change in behavior; hence it offers

more realistic predictions of new cases over time and also performs well even with the limited data available during the epidemic's initial phases.

In this paper we develop and compare two novel approaches to better allocate epidemic intervention resources across regions. We calibrate an epidemic model to the 2014 Ebola outbreak, (Section 4.3), although the framework could apply to any emerging infectious disease outbreak. The first allocation approach, *Greedy* R_0, simply rank-orders affected regions based on the basic reproduction number, R_0, and allocates resources only to regions where $R_0 > 1$, which indicates a growing epidemic (Section 4.4). The second approach, *ADP algorithm*, formulates the resource allocation problem as a nonlinear optimization problem with an underlying dynamic epidemic model, which we numerically solve using an Approximate Dynamic Programming (ADP) algorithm. While the *Greedy* R_0 approach is relatively easy to compute, the *ADP algorithm* provides a better solution when more complex allocation decisions are permitted and the impact on future transmission is included. Using Ebola incidence data from the 2014 epidemic in West Africa, we consider both approaches under different resource budgets and epidemic forecasting scenarios, and compare their relative performance to the planned allocation and various heuristics (Section 4.5). Finally, we discuss implications for practice and key conclusions (Section 4.6).

4.2 Literature Review

4.2.1 Epidemic Modeling of Ebola

Prior studies have developed models of human-to-human transmission of major Ebola outbreaks; for animal-to-human transmission, see Walsh et al. (2005); Groseth et al. (2007). Most studies forecast the epidemic's trajectory using an extension of a compartmental susceptible-infectious-removed (SIR) model, which specifies a system of nonlinear differential equations to model the movement of individuals between compartments (Anderson & May, 1991). In such models, **S**usceptible individuals are vulnerable to infection, **I**nfectious individuals can transmit Ebola to those susceptible, and **R**emoved individuals have either died or recovered from the disease.

A key measure of an epidemic's persistence in a population is the *basic reproduction number*, R_0, defined as the average number of secondary infections caused by a typi-

cal infected individual in a predominantly susceptible population. If $R_0 < 1$ then the epidemic eventually dies out, but if $R_0 > 1$ then the disease remains endemic. R_0 is essentially the sufficient contact rate for transmission multiplied by the duration of infectivity. Therefore, R_0 could be lowered by (i) reducing contact between susceptible and infectious people (e.g., through quarantines or closures of school and work), (ii) reducing the likelihood of transmission if contact does occur (e.g., through bleaching stations or personal protective equipment for health care workers), or (iii) reducing the period of infectivity (e.g., by increasing medical personnel and treatment centers or ensuring safe burial practices).

The earliest Ebola outbreaks to be extensively modeled were the 1995 Democratic Republic of the Congo (DRC) and 2000 Uganda epidemics (Table 4.1). Chowell et al. (2004) used a susceptible-exposed-infectious-removed (SEIR) model with least-squares to estimate R_0 values of 1.34 for Uganda and 1.83 for DRC, which likely differ because of different virus subtypes. Using a similar SEIR model, Lekone & Finkenstädt (2006) estimated an R_0 of 1.36 for the 1995 DRC outbreak. Another modeling study (Legrand et al., 2007) of the DRC outbreak concluded that a delay in intervention deployment was the most important factor in determining final epidemic size, followed by speed of hospitalization and efficacy of burial interventions. Ferrari et al. (2005) estimated R_0 while correcting for underreporting and the bias associated with using discrete data within a continuous process; their R_0 estimates were generally higher than those of other studies for these particular epidemics.

Recent models of the 2014 West Africa Ebola outbreak have estimated an R_0 of approximately 1.80 (range: 1.51 to 2.53), with every study concurring that $R_0 > 1$ (Table 4.1). Different assumptions regarding intervention efficacy, time horizons, underlying transmission, or mortality rates—combined with a lack of high-quality case data—contribute to the differing estimates of R_0. Most studies calibrate their models with country-level case counts to forecast the epidemic's trajectory. Using data from Sierra Leone and Liberia, Meltzer et al. (2014) found that rapid hospitalization of Ebola-stricken patients could reduce transmission risk and help contain the epidemic. The World Health Organization (WHO Ebola Response Team, 2014) predicted future cases by assuming a Poisson process based on field data; this analysis incorporated transmission correlations between districts but did not explicitly model dynamic transmission. Althaus (2014) used a prior model (Legrand et al., 2007) and found that greater intervention efficacy slowed outbreaks in Guinea and Sierra

Table 4.1: Estimated R_0 values for major Ebola outbreaks.

Year	Region	R_0	Source
1995	Democratic Republic of the Congo	1.83	Chowell et al. (2004)
		3.65	Ferrari et al. (2005)
		1.36	Lekone & Finkenstädt (2006)
2000	Uganda	1.34	Chowell et al. (2004)
		1.79	Ferrari et al. (2005)
2014	Guinea	1.51	Althaus (2014)
		1.71	WHO Ebola Response Team (2014)
	Liberia	1.59	Althaus (2014)
		2.22	Rivers et al. (2014)
		1.83	WHO Ebola Response Team (2014)
		1.84	Merler et al. (2015)
	Montserrado County, Liberia	2.49	Lewnard et al. (2014)
	Sierra Leone	2.53	Althaus (2014)
		1.78	Rivers et al. (2014)
		2.02	WHO Ebola Response Team (2014)
	Liberia and Sierra Leone	1.80	Meltzer et al. (2014)
	All affected regions	1.78	Fisman et al. (2014)
		1.80	Gomes et al. (2014)

Leone whereas transmission remained constant in Liberia. Khan et al. (2015) used an extended compartmental model and estimated R_0 values of 1.76 for Liberia and 1.49 for Sierra Leone. Using a compartmental model and a Markov Chain Monte Carlo (MCMC) approach for fitting data from Montserrado County, Liberia, Lewnard et al. (2014) found that the effectiveness of Ebola treatment units depends on how rapidly cases are detected.

Several studies have proposed models to account for time-dependent intervention efficacy or heterogeneous populations. Fisman et al. (2014) combined R_0 with a decay factor representing the net effect of control measures against further transmission; however, the authors modeled the epidemic by aggregating populations at the country level. Merler et al. (2015) also used a model based on Legrand et al. (2007) and estimated an R_0 of 1.84 for Liberia. Instead of assuming homogeneous mixing, these authors divided the population into households and derived a mixing pattern from the behavior of individuals in households. They estimated that 52.9% of the infections occur in households or in the general community. Other epidemic models with multiple dependent populations have been applied to influenza (Wu et al., 2007) as well as other disease outbreaks (Apolloni et al., 2014), but these models have not been

applied to Ebola.

4.2.2 Resource Allocation for Infectious Diseases

Optimization models of resource allocation for other infectious diseases typically consider independent populations or just a few dependent populations (Greenhalgh, 1986; Wilson et al., 2006; Rowthorn et al., 2009). Given Ebola's long incubation period—up to 21 days after infection before the onset of symptoms (World Health Organization, 2015)—exposed individuals may move between regions and interact with different populations before becoming aware of their infection status. Especially when resources are constrained, policy makers and public health leaders must balance the trade-off between prioritizing highly affected key populations and providing a solution that is equitable; this trade-off was exemplified during the 2009 H1N1 influenza pandemic, when vaccines were in short supply (Medlock & Galvani, 2009). Reaching agreement on a stated objective, such as minimizing new infections or maximizing life-years, poses an additional challenge. Brandeau (2004) reviewed the literature on optimal resource allocation in epidemic settings.

An early study (Greenhalgh, 1986) used control theory to determine the optimal rate at which people within a single population should be removed. Optimal control has also been used in the more complex case of multiple interacting populations. Rowthorn et al. (2009) used a susceptible-infectious-susceptible (SIS) model to allocate a treatment budget to two connected populations; they found that onward transmission was minimized by first treating the population with *fewer* infectious individuals. Conversely, Ndeffo Mbah & Gilligan (2011) used optimal control with an SIR model with ten populations and found that vaccinating populations with *more* infectious individuals was optimal. Richter et al. (1999) drew similar conclusions; they found that prioritizing HIV prevention efforts among high-risk populations minimized new infections over a five-year period.

Spatial differences between populations have been considered within compartmental epidemic models. Wilson et al. (2006) compared the allocation of anti-retroviral therapy among urban and rural areas, and Wu et al. (2007) used a nonlinear optimization program to evaluate the allocation of influenza vaccinations among ten regions in the United States. Gerberry et al. (2014) used a linear program to maximize HIV in-

fections averted with different prevention plans in nine South African regions; these authors found that resources should be provided to the most susceptible individuals. By constructing nonlinear production functions that relate the investment in intervention to its efficacy, several studies have optimized resource allocation in dependent populations (Zaric & Brandeau, 2001, 2002) or in multiple independent populations (Brandeau et al., 2003).

Despite numerous studies in the epidemiology and operations research literature, there is no comprehensive model for efficient resource allocation during a rapidly evolving outbreak. Although spatial heterogeneity has been previously examined (Wilson et al., 2006; Wu et al., 2007), those studies considered neither the epidemic's future trajectory nor the behavioral responses of individuals to epidemic conditions. We therefore contribute to the literature in two ways. We first develop a multi-population epidemic model with a novel behavioral transmission dampening factor. We secondly develop and test two resource allocation approaches: a *Greedy* R_0 approach based on an analytically derived bound on R_0, which could be used to make a fast decision when the allocation time frame is short, and an *ADP algorithm*, which could provide an improved solution when the allocation decision is more complex, e.g., given a larger budget to allocate, or constraints on the speed of deployment. Both approaches incorporate spatial heterogeneity among populations, and thereby support informed decision making in particular during an epidemic's early stages—when accurate and timely case data are limited.

4.3 Epidemic Model

We construct a dynamic epidemic model with a mixing coefficient to capture Ebola transmission both within a geographic region and between regions (Figure 4.1). We introduce a novel, time-dependent transmission dampening coefficient based on the degree of behavioral adaptation of the affected population as existence of the outbreak becomes apparent. Parameter values are estimated numerically by fitting the model to daily case counts for the 2014 Ebola epidemic in West Africa.

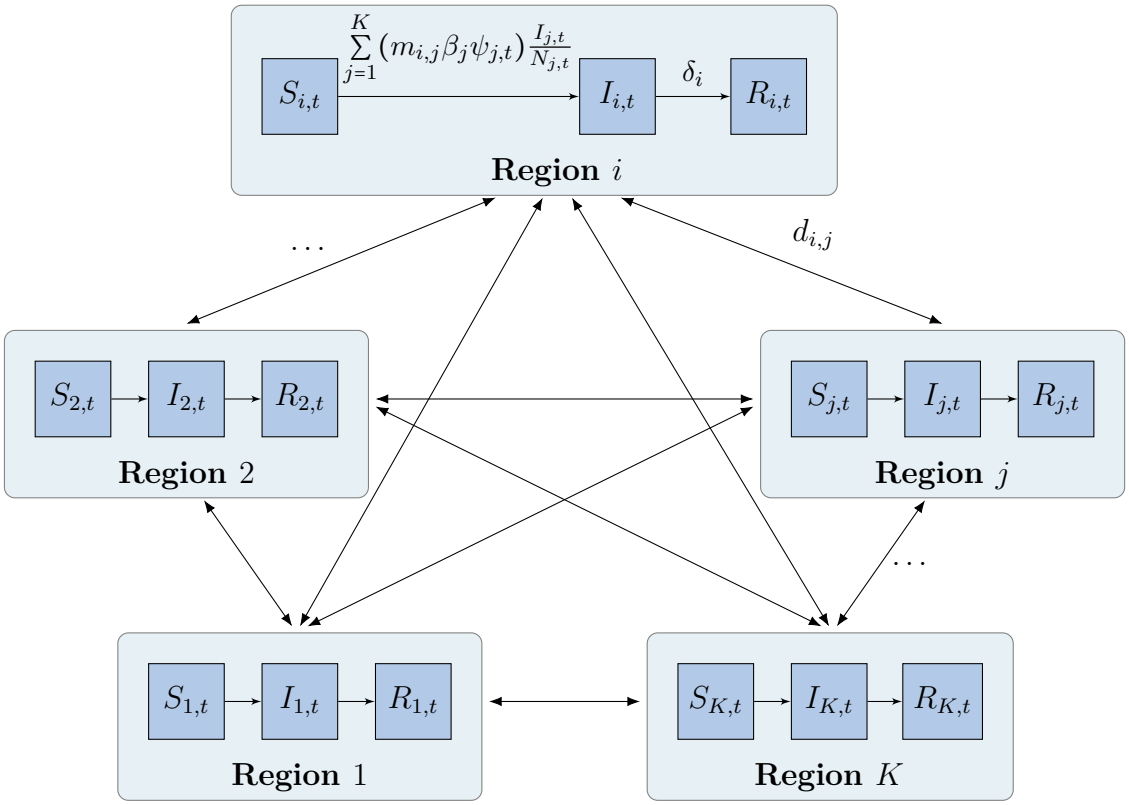

Figure 4.1: Schematic diagram of Susceptible-Infectious-Removed (SIR) epidemic model among K interdependent regions. Dark blue boxes depict Ebola infection status, light blue boxes depict geographic regions, and arrows between regions show Ebola transmission rates inversely proportional to geographic distance.

4.3.1 Infection Dynamics

To model Ebola transmission from infected to uninfected individuals, as well as recovery or death, we assume an SIR compartmental model for each of K geographic regions. The epidemic's nonlinear dynamics arise because the rate of new infections is proportional to the sizes of the susceptible population in region i and the infectious population in region j.

Let $S_{i,t}$, $I_{i,t}$, and $R_{i,t}$ denote the number of susceptible, infectious, and removed individuals, respectively, in region i at time t. $N_{i,t}$ denotes the total population, such that $S_{i,t} + I_{i,t} + R_{i,t} = N_{i,t}$. The *transmission coefficient*, β_j, is the per-contact transmission probability between a susceptible person and an infectious person from region j. If transmission occurs, a newly infected individual can then transmit the infection after

symptom onset until recovery or burial, for a period of time with mean duration $1/\delta_i$. The full model is represented by Equations (4.3)–(4.5).

Most epidemic models assume *homogeneous mixing* within the population, which implies that susceptible individuals randomly come into contact with other individuals in the population with equal probability. This assumption is reasonable for individuals living near each other, but it may not hold for geographically distant populations. One solution is to assume that populations are completely isolated, so that transmission occurs only inside each geographic region; thus $\frac{dS_{i,t}}{dt} = -S_{i,t}\beta_i\frac{I_{i,t}}{N_{i,t}}$, $\forall i$. Although this approach simplifies the dynamics, it may not adequately reflect the natural spread of an infectious disease like Ebola.

4.3.2 Geographic Proximity

To forecast such a large-scale and heterogeneous epidemic, we extend the basic SIR model to capture interactions among individuals from different regions by adjusting the effective transmission rate to reflect geographic proximity of populations to one another. Let c_0 $(0 \leq c_0 \leq 1)$ denote the fraction of contacts located within the home region, with the remaining $1-c_0$ in outside regions. The degree of mixing $m_{i,j}$ between individuals in regions i and j is proportional to the *inverse* of the distance $d_{i,j}$ between the two regions' capitals (Figure 4.1), adapting the approach of Wu et al. (2007):

$$m_{i,j} = \begin{cases} (1-c_0)\dfrac{\frac{1}{d_{i,j}}}{\sum\limits_{\substack{k=1 \\ k\neq j}}^{K} \frac{1}{d_{k,j}}} & \text{if } i \neq j, \\ c_0 & \text{if } i = j. \end{cases} \tag{4.1}$$

4.3.3 Dampening via Adaptive Behavior

During the outbreak's early phases, Ebola case numbers increased by 91% within 24 days starting in August 2014. Yet over an equal period in March 2015, cases increased by only 3% despite a large fraction of the population remaining susceptible (Humanitarian Data Exchange, 2015a). Although intervention measures were simultaneously scaled up in the intervening months, disease transmission was reduced further owing to behavioral changes made by individuals living in Ebola-afflicted communities (Del

Valle et al., 2005). We model this change in disease contagion by introducing the *dampening coefficient* $\psi_{j,t}$, which reduces the effective transmission rate in region j as the epidemic progresses and individuals modify their behavior accordingly. We assume a logistic function that allows for a slow increase in awareness during the outbreak's beginning (as captured by the steepness factor α_j for each region), followed by rapid behavioral adaptation as the epidemic peaks until the minimum dampening value $1 - \psi_j^{\min}$ for region j is reached:

$$\psi_{j,t} = 1 - \frac{\psi_j^{\min}}{1 + e^{-\alpha_j (t - T_j/2)}}. \tag{4.2}$$

The period of epidemic data for region j used for calibrating the model is T_j, with $T_j/2$ representing the midpoint. For each region j, we numerically estimate the parameters α_j and $\psi_j^{\min}$.

We combine an SIR model with the geographic mixing component ($m_{i,j}$) and behavior dampening coefficient ($\psi_{j,t}$) to obtain the following system of equations for region $i = 1, 2, \ldots, K$:

$$\frac{dS_{i,t}}{dt} = -S_{i,t} \sum_{j=1}^{K} (m_{i,j} \beta_j \psi_{j,t}) \frac{I_{j,t}}{N_{j,t}}, \tag{4.3}$$

$$\frac{dI_{i,t}}{dt} = S_{i,t} \sum_{j=1}^{K} (m_{i,j} \beta_j \psi_{j,t}) \frac{I_{j,t}}{N_{j,t}} - \delta_i I_{i,t}, \tag{4.4}$$

$$\frac{dR_{i,t}}{dt} = \delta_i I_{i,t}. \tag{4.5}$$

The term $(m_{i,j} \beta_j \psi_{j,t})$ is the effective transmission rate at time t from infectious individuals in region j to susceptible individuals in region i. We conduct sensitivity analysis by ignoring geographic distance ($c_0 = 1$) or behavior dampening ($\psi_{i,t} = 1$, $\forall i$, $0 \le t \le T_i$); results are presented in Section 4.5.1.

The epidemic model was calibrated with confirmed Ebola case data from 21 regions in Guinea, Liberia, and Sierra Leone, which collectively accounted for 60% of the total infections observed (Humanitarian Data Exchange, 2015b). We used between 20 and 100 days of the available 140 days of data to test the model's performance against the observed case count. Model calibration was performed using a Markov chain Monte Carlo (MCMC) approach with a Metropolis–Hastings algorithm, an established method for fitting epidemic models to empirical data (Gibson, 1998; Currie, 2007). A

detailed description of the calibration approach is available in the Appendix.

4.3.4 R_0 Computation

To compute the basic reproduction number, R_0, of our SIR model for K regions with transmission dependent on geographic proximity and dynamic behavior dampening, we use the *next generation matrix* approach (van den Driessche & Watmough, 2002; Arino & Van Den Driessche, Pauline, 2003). This method essentially linearizes the nonlinear system (4.3)–(4.5) near the disease-free equilibrium, $I_i(t) = 0, \forall i$, to determine if the epidemic dies out or remains endemic in the overall population.

We first introduce the following vector notation (the subscript t is suppressed for clarity). Let $\mathbf{s} = (S_1, \ldots, S_K)'$, $\mathbf{i} = (I_1, \ldots, I_K)'$, $\mathbf{r} = (R_1, \ldots, R_K)'$, $\mathbf{n} = (N_1, \ldots, N_K)'$, and $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_K)'$. Let $diag(\mathbf{s})$ be a $K \times K$ matrix with diagonal components of the vector $\mathbf{s}$. Define the $K \times K$ matrix $\mathbf{A}$ with $A_{i,j} = m_{i,j}\beta_j\psi_j/N_j$. The epidemic model can be rewritten as:

$$\dot{\mathbf{s}} = -diag(\mathbf{s})\mathbf{A}\mathbf{i}, \tag{4.6}$$

$$\dot{\mathbf{i}} = diag(\mathbf{s})\mathbf{A}\mathbf{i} - diag(\boldsymbol{\delta})\mathbf{i}, \tag{4.7}$$

$$\dot{\mathbf{r}} = diag(\boldsymbol{\delta})\mathbf{i}. \tag{4.8}$$

with $\mathbf{s} + \mathbf{i} + \mathbf{r} = \mathbf{n}$.

Definition 4.1. *A square non-negative matrix* $\mathbf{F}$ *is called primitive if there is an* n *such that all the components of* $\mathbf{F}^n$ *are positive.*

Assumption 4.1. *Infected individuals must eventually recover or die, so* $\delta_i > 0, \forall i$.

Assumption 4.2. *After a sufficient number of time periods, an infection originating in region* j *will ultimately lead to an infection in region* i*; therefore,* $\mathbf{A}$ *is primitive.*

Lemma 4.1. *Given a primitive matrix* $\mathbf{F}$ *and positive matrix* $\mathbf{V}$*, then* $\mathbf{F}\mathbf{V}^{-1}$ *has a positive real eigenvalue* λ_{max} *such that all other eigenvalues of* $\mathbf{F}\mathbf{V}^{-1}$ *satisfy* $|\lambda| \leq \lambda_{max}$.

Proof The proof of Lemma 4.1 follows from the standard Perron-Frobenius theorem for non-negative matrices.

Using the next generation matrix approach, the basic reproduction number for our model is:

$$R_0 = \lambda_{\max}(\mathbf{F}\mathbf{V}^{-1}) = \lambda_{max}\left(diag(\mathbf{s})\mathbf{A}diag(\boldsymbol{\delta}^{-1})\right), \tag{4.9}$$

where $\lambda_{\max}(\cdot)$ is the dominant eigenvalue of a matrix. Component i, j of the matrix $\mathbf{F} = diag(\mathbf{s})\mathbf{A}$ is the rate of new infections occurring in region i that originate in region j. Component j, k of the matrix $\mathbf{V}^{-1} = diag(\boldsymbol{\delta}^{-1})$ is the average length of time that an individual in region j is infectious before recovery or death. Hence, component i, k of $\mathbf{F}\mathbf{V}^{-1}$ is the expected number of people in region i who acquire infection from an infectious individual originally located in region k at the epidemic's start.

Theorem 4.1. *Given a deterministic epidemic model in multiple dependent populations with transmission matrix* $\mathbf{A}$ *and recovery or death rate* $\boldsymbol{\delta}$*, then the epidemic persists if* $R_0 > 1$ *and dies out if* $R_0 < 1$*.*

Proof With no population re-entry or new births, the number of susceptible individuals in each region is monotonically decreasing, so $\dot{\mathbf{s}} \leq 0$. Because the system of equations (4.6)–(4.8) is Lipschitz continuous, the nonlinear differential equation (4.7) can be approximated by the linear differential equation:

$$\dot{\mathbf{x}} = diag(\mathbf{s})\mathbf{A} - diag(\boldsymbol{\delta})\mathbf{x}$$

at points near the epidemic's start. The solution of this is of the form:

$$x_i(t) = \sum_{j=1}^{K} c_{i,j} e^{-\lambda_j t}$$

with constants $c_{i,j}$. The trajectory of $\mathbf{x}$ will approach the disease-free equilibrium $x_i(t) = 0, \forall i$, if the maximum real part of all the eigenvalues of the matrix $diag(\mathbf{s})\mathbf{A} - diag(\boldsymbol{\delta})$ does not exceed 0. Finally,

$$\lambda_{max}\left(diag(\mathbf{s})\mathbf{A} - diag(\boldsymbol{\delta})\right) < 0 \Longleftrightarrow \lambda_{max}\left(diag(\mathbf{s})\mathbf{A}diag(\boldsymbol{\delta}^{-1})\right) < 1$$

Therefore, the epidemic will approach the disease-free equilibrium if $R_0 < 1$.

Corollary 4.1. *The dominant eigenvalue* λ_{max} *has the following bounds:*

$$\min_j \sum_{i=1}^{K} \frac{m_{i,j}\beta_j\psi_j}{\delta_j} \leq \lambda_{max} \leq \max_j \sum_{i=1}^{K} \frac{m_{i,j}\beta_j\psi_j}{\delta_j},$$

$$\min_i \sum_{j=1}^{K} \frac{m_{i,j}\beta_j\psi_j}{\delta_j} \leq \lambda_{max} \leq \max_i \sum_{j=1}^{K} \frac{m_{i,j}\beta_j\psi_j}{\delta_j}.$$

Proof The proof of Corollary 4.1 follows from the standard Perron-Frobenius theorem for non-negative matrices.

Using Corollary 4.1 for the epidemic model in Equations (4.3)–(4.5), we can provide upper bounds for R_{0j}, the basic reproduction number in each region $j = 1, \ldots, K$. R_{0j} is essentially the total number of new infections—that occur both within region j and to all outside regions—caused by each infectious individual originally in region j at the epidemic's start:

$$R_{0j} \leq \sum_{i=1}^{K} \frac{m_{i,j}\beta_j\psi_j}{\delta_j}. \tag{4.10}$$

4.4 Allocation of Intervention Resources

During epidemic crises, policy makers and non-government organizations (NGOs) must often allocate resources based on the limited data available at the time of decision (Doctors Without Borders, 2015). Our calibrated epidemic model can help improve these decisions by augmenting existing data with projections of the epidemic's trajectory over time and across regions. Identifying the optimal distribution of resources will maximize their effect. An overview of the model elements and their interdependencies is presented in Figure C.1 in the Appendix.

Compared to transmissible diseases with long infectious periods, such as HIV/AIDS, patients infected with Ebola typically suffer a relatively sudden death or recovery. Maintaining sufficient ETU capacity is key to any containment strategy (Lewnard et al., 2014; Merler et al., 2015), requiring both physical bed infrastructure and health care workers. Infected individuals who are treated in ETUs are less likely to transmit the virus, owing to better protective measures by health care workers and reduced contact with susceptible individuals in the community. We therefore assume that the transmission coefficient β_j of those treated within an ETU is reduced by a factor ξ ($0 \leq \xi \leq 1$). Although we consider ETU beds as our primary resource, the optimization model could apply to other resources, including health care personnel, vaccinations, and medical equipment. The remainder of this section describes our two resource

allocation approaches.

4.4.1 Greedy R_0

By Theorem 1, we ensure that the epidemic theoretically dies out if $\max_j R_{0j} < 1$. Using Equation (4.10), we use the upper bound of R_{0j} for each region as a proxy for the infection potential of that region. Consider a resource allocation scheme that rank-orders regions based on R_{0j}:

$$R_0^{(1)} \leq R_0^{(2)} \leq \cdots \leq R_0^{(K-1)} \leq R_0^{(K)}$$

where $R_0^{(i)}$ is the i^{th}-smallest value. In each period, r resources are distributed (the maximum deployment rate) to any region j where $R_{0j} > 1$, in descending order of R_{0j} until the budget is exhausted. The resulting allocation heuristic therefore prioritizes regions where an epidemic is progressing, but ignores regions where the infection potential is not yet high enough or the epidemic has already abated.

4.4.2 ADP Algorithm

Nonlinear Optimization Model

We also formulate a nonlinear program with underlying dynamics governed by the model in Equations (4.3)–(4.5). We divide the time horizon into discrete, equidistant periods $t \in (t_0, \ldots, t_f)$; in the Ebola example, each period consists of one week. In period t, region i is assigned $a_{i,t}$ resources. The allocation is constrained by the total available budget B, and the maximum deployment rate per region per period r, reflecting the limited ability to distribute available resources quickly in affected regions. Once resources are assigned to a region, we assume that they remain there for the rest of the outbreak: $a_{i,t} \geq a_{i,t-1}\ \forall i, t > 0$. This assumption is reasonable in the context of Ebola, although it could be relaxed. In our example, $a_{i,t}$ denotes an integer number of beds; however, this approach could be generalized to noninteger variables.

The projected number of new Ebola infections $I_{i,t}^{\text{new}}$ in region i during period t is based

on the underlying SIR model:

$$I_{i,t}^{\text{new}}(a_t) = S_{i,t} \sum_{j=1}^{K} (m_{i,j}\beta_j\psi_{j,t}) \left(\frac{\xi \min(I_{j,t}, a_{j,t}) + (I_{j,t} - a_{j,t})^+}{N_{j,t}} \right). \tag{4.11}$$

The number of infected individuals treated in ETU beds in region j during period t is $\min(I_{j,t}, a_{j,t})$, where the effective transmission rate of this population is reduced by ξ. The term $(I_{j,t} - a_{j,t})^+$ corresponds to any remaining infected individuals who are not treated in an ETU—because of a bed shortage, for example—and who therefore continue to transmit infection at the original rate.

Our objective is to minimize the cumulative number of new Ebola infections in all regions over the entire time horizon. One might instead consider other objectives, such as maximizing QALYs or minimizing the number of infections in a given region. Under the assumptions given above, we obtain the following optimization model:

$$\begin{aligned}
\text{minimize} \quad & \sum_{i=1}^{K} \sum_{t=t_0}^{t_f} I_{i,t}^{\text{new}}(a_t) && (4.12)\\
\text{subject to} \quad & \sum_{i=1}^{K} a_{i,t} \le B \quad \forall t, && (4.13)\\
& 0 \le (a_{i,t} - a_{i,t-1}) \le r \quad \forall i, t > 0, && (4.14)\\
& a_{i,t} \text{ integer} \quad \forall i, t, && (4.15)\\
& \text{equations (4.3)–(4.5).} && (4.16)
\end{aligned}$$

Dynamic Program

To solve this nonlinear program, we use deterministic DP with a finite horizon, extending the approach of Blount et al. (1997) by optimizing resource allocation across multiple dependent populations.

In period t, the state space x_t consists of two vectors $(\mathbf{S}_t,\ \mathbf{I}_t)$ consisting of K elements each:

$$\mathbf{S}_t = (S_{1,t}, S_{2,t}, \ldots, S_{K,t}) \quad \text{and} \quad \mathbf{I}_t = (I_{1,t}, I_{2,t}, \ldots, I_{K,t}).$$

Here $\mathbf{S}_t$ and $\mathbf{I}_t$ represent the number of susceptible and infectious individuals, respectively, in each region during period t. State transitions from x_t to x_{t+1} are governed not only by the number of susceptible and infectious individuals but also by the resources

allocated during period t:

$$S_{i,t+1} = S_{i,t} - I^{\text{new}}_{i,t}(x_t, a_t), \tag{4.17}$$

$$I_{i,t+1} = I_{i,t} + I^{\text{new}}_{i,t}(x_t, a_t) - \delta_i I_{i,t}. \tag{4.18}$$

The initial conditions, $x_0 = (\mathbf{S}_{t_0},\ \mathbf{I}_{t_0})$, refer to the size of the susceptible and infectious populations in each region at the start of the epidemic. We assume that no ETU resources are initially available within each region ($a_{i,t_0} = 0,\ \forall i$), which is reasonable for an emerging outbreak like Ebola, although this assumption could be relaxed. This DP problem can be divided into a sequence of subproblems and then solved, by the Principle of Optimality (Bertsekas, 1995), with a terminal value $V_{t_f}(x_{t_f}) = 0$. The solution is obtained via backward recursion starting at t_f while assuming that no future infections occur after this period—in other words, that the outbreak eventually dies out.

Approximate Dynamic Program

With 21 regions, home to several million susceptible individuals, finding the exact solution of this dynamic program is computationally intractable given the "curse of dimensionality" caused by the large state space. To overcome this issue, we propose an ADP approach that approximates the value of a state without sweeping over the entire state space. This methodology has previously been applied to resource allocation problems in health care, such as planning elective patient admissions (Barz & Rajaram, 2015; Hulshof et al., 2015), but not to epidemic control.

To demonstrate how our solution method could apply to other intractable dynamic programs, we proceed as follows. First we introduce the policy iteration algorithm used and then reformulate each period's decision into multiple Knapsack subproblems; this allows ADP to be used with large state spaces. Next we explain the selection and use of linear basis functions to approximate the value of a state, and finally we describe the method used to explore or exploit the state space. The psuedocode for the developed ADP is presented in Fig. C.2 in the Appendix.

Policy Iteration.

Our ADP approach is based on a policy iteration algorithm (Powell, 2011), which solves a finite-horizon problem using linear basis functions to approximate the value function. We replace the exact value of the state $V_t(x_t)$ with a value approximation $\hat{V}_t(x_t)$:

$$\hat{V}_t(x_t) = \min_{a_t \in \mathcal{A}_t} \left(\sum_{i=1}^{K} I_{i,t}^{\text{new}}(x_t, a_t) + \gamma \hat{V}_{t+1}(x_{t+1}) \right). \tag{4.19}$$

The ADP algorithm (see Appendix for details) consists of one outer loop and two nested loops. The first nested loop computes an optimal allocation decision, based on the current approximate value function using a nonlinear Knapsack approach presented in Section C.2.1 in the Appendix. The second nested loop updates the basis function parameter vector ($\boldsymbol{\theta}^n$) with the observed state value. The outer loop repeats the two inner loops for N iterations, where N is sufficiently large to reach convergence of $\boldsymbol{\theta}$ in all tested cases. The optimal policy can then be computed using $\boldsymbol{\theta}^N$ within the value function approximation.

Feature Extraction.

To efficiently approximate the future value of a state x_t, without sweeping over the entire state space, we use a basis function approach (Powell, 2011). A basis function $\phi_f(x_t)$ summarizes just a few state features $f \in \mathcal{F}$, thereby reducing the large set of state variables from the original dynamic program. For our Ebola model, we examine the following set of features:

$$\mathcal{F} = \left\{ S_i,\ I_i,\ R_0,\ S_i I_i,\ \frac{I_i}{S_i},\ \sum_{j=1}^{K} I_j,\ S_i \sum_{j=1}^{K} I_j,\ \sum_{j=1}^{K} \frac{I_j}{S_i},\ R_{0i} S_i I_i,\ B \quad \forall i \right\}. \tag{4.20}$$

Selecting appropriate basis functions is essential for a good approximation of the value function because doing so extracts those elements from the state space that have high explanatory power (e.g., the number of infected individuals in a population affects the future epidemic trajectory in that population). We rank potential basis functions in terms of their explanatory power for a state's future value (Hulshof et al., 2015). We simulate epidemic trajectory paths under different resource allocation policies, and then treat these simulated values as observations in a linear regression to determine

Table 4.2: Basis functions considered and the average R^2 and Akaike information criterion (AIC) values across all populations.

Basis function ($\forall i \in 1, \ldots, K$)	# of variables	R^2	AIC ($\times 10^5$)
$S_i I_i$	K	.653	2.42
$I_i + S_i \sum_{j=1}^{K} I_j$	$2K$	.653	2.43
$S_i \sum_{j=1}^{K} I_j + B$	$K+1$	.771	2.36
$\frac{I_i}{S_i} + B$	$K+1$	.822	2.30
R_{0i}	K	.592	2.49
$R_{0i} S_i I_i + B$	$K+1$	.867	2.09

which set of features best approximates the value function (Table 4.2).

We use $\boldsymbol{\phi}(x_t)$ to denote the column vector of features and use $\boldsymbol{\theta}_t$ to denote the parameter vector of estimated regression coefficients. The value function approximation at iteration n of the ADP algorithm is:

$$\bar{V}_t^n(x_t|\boldsymbol{\theta}_t^n) = \boldsymbol{\phi}(x_t)^T \boldsymbol{\theta}_t^n. \tag{4.21}$$

$$\lambda^n = 1 - \frac{\mu}{n}. \tag{4.22}$$

To assure that the ADP is not only visiting states with an already high approximated value we use ϵ-greedy exploration throughout our experiments. Detailed reasons for the selection of this method are given in Section C.2.3 in the Appendix. Let $\boldsymbol{\phi}(x_t)$ denote the column vector of features, and $\boldsymbol{\theta}_t$ denote the parameter vector of estimated regression coefficients. The value function approximation at iteration n of the algorithm is:

$$\bar{V}_t^n(x_t|\boldsymbol{\theta}_t^n) = \boldsymbol{\phi}(x_t)^T \boldsymbol{\theta}_t^n. \tag{4.23}$$

To ensure that the algorithm is not only visiting states with an already high approximated value, we use ϵ-greedy exploration throughout our experiments. The algorithm for updating θ_n^t after each iteration is presented in Section C.2.2.

4.5 Numerical Results

We first describe fitting our SIR model with behavior dampening and proximity-based transmission to the 2014 Ebola epidemic in West Africa. To illustrate the utility of the

Greedy R_0 approach and *ADP algorithm*, we compare their performance with other benchmark heuristics. Calculations were performed in Matlab R2015b.

4.5.1 Epidemic Model Fit

Ebola case count data for all 21 regions were obtained from the Humanitarian Data Exchange (2015b); total population estimates N_i for each region were obtained from Sierra Leone Statistics (2004); Liberia Institute of Statistics and Geo-Information Service (2008); Institut National de la Statistique (2014). The mean time from symptom onset to recovery or burial of an individual, $1/\delta_i$, was based on prior estimates from Merler et al. (2015). The reduction in infectivity among individuals treated in ETUs was set to $\xi = 0.5$, the most conservative value provided by Merler and colleagues. An MCMC approach was used to numerically estimate remaining parameters: the transmission coefficient β_i, minimum dampening value $\psi_i^{\min}$, and steepness factor α_i for each region as well as the proportion c_0 of contacts in each region. MCMC was run for 200,000 iterations, though our final parameter estimates were based on only the last 50,000 iterations. Details on the model calibration as well as convergence graphs are presented in Section C.3 in the Appendix.

Using the method in Section 4.3.4 and parameter values estimated using 100 days of data, we obtain bounds for the basic reproduction number of $0.19 \leq R_0 \leq 1.69$, and a precise estimation through Equation (4.9) of $R_0 = 1.17$. This estimate is significantly lower than earlier published estimates (Table 4.1), as our model accounts for a reduction in transmission through behavior change.

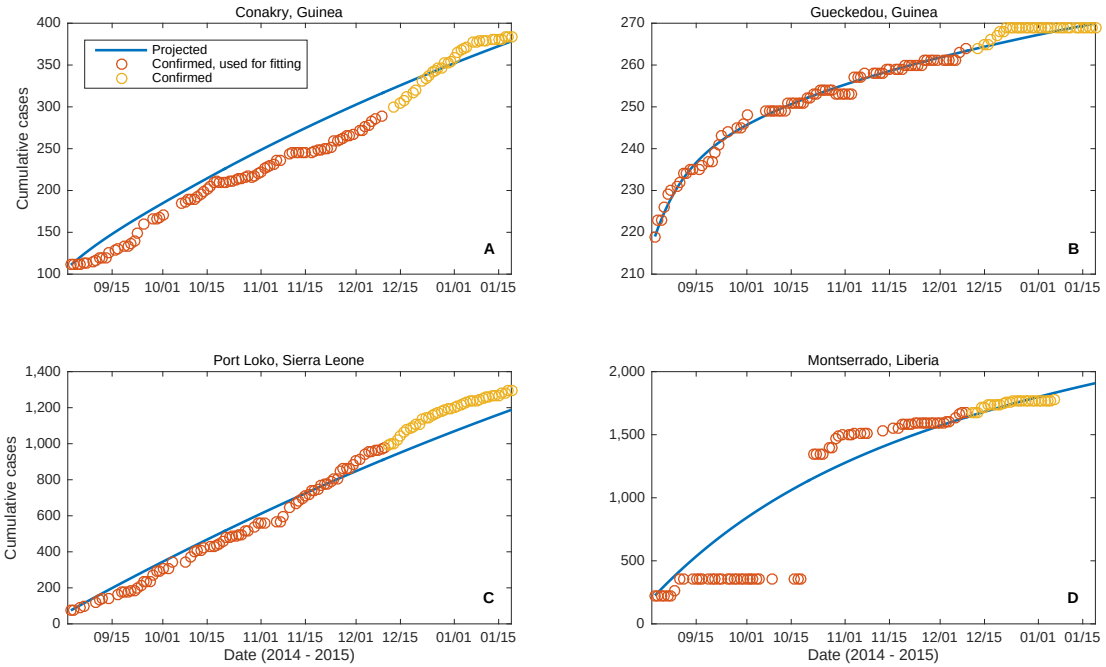

Figure 4.2: Observed and projected Ebola cases for (A) Conakry, Guinea, (B) Guéckédou, Guinea, (C) Port Loko, Sierra Leone, and (D) Montserrado, Liberia; the model is fit to 100 days of data (red circles) and compares the following 40 days of actual cases (yellow circles) with projected cases (blue line).

After estimating all model parameters, we forecast new Ebola cases over a 40-day period (Table 4.3). Using 100 days of past data, the model projects 9,555 infections across all regions over the following 40 days, compared to 9,480 confirmed cases during this period—an aggregate overestimate of less than 1%. The mean absolute percentage error (MAPE) across all regions is 11% (Figure 4.2), which may reflect underreporting and/or variations in intervention rollout. Estimating parameters based on 60 days of past data increases the MAPE to 20%, and with only 20 days of past data, the MAPE reaches 30% (Table 4.3). These figures demonstrate that our model yields accurate forecasts when sufficient data are available yet still provides a good approximation of the epidemic's future trajectory when data are limited. Although our model's validity deteriorates with fewer data points, it performs far better than the initial projections of the 2014 Ebola epidemic. If we assume static behavior dampening coefficients $\psi_i(t)$, then the MAPE increases from 11% to 24% using 100 days of data, and up to 133% with only 20 days of data (Table 4.3). Detailed epidemic model fit results for all regions are presented in Section C.3.1.

Table 4.3: Observed and projected Ebola cases with mean absolute percentage error (MAPE) averaged across 21 regions for different model configurations; projected cases are compared with actual cases observed 40 days after the calibration period.

Fitting period	Observed cases	Base model		No behavior damp.		No connectivity	
		Projected cases	MAPE (%)	Projected cases	MAPE (%)	Projected cases	MAPE (%)
20	6,361	3,823	30	1,228	166	3,845	35
40	7,527	4,705	27	5,093	24	6,839	50
60	8,468	7,006	20	8,814	23	12,818	63
80	9,206	8,641	15	9,792	15	16,232	70
100	9,480	9,555	11	13,024	24	16,991	68

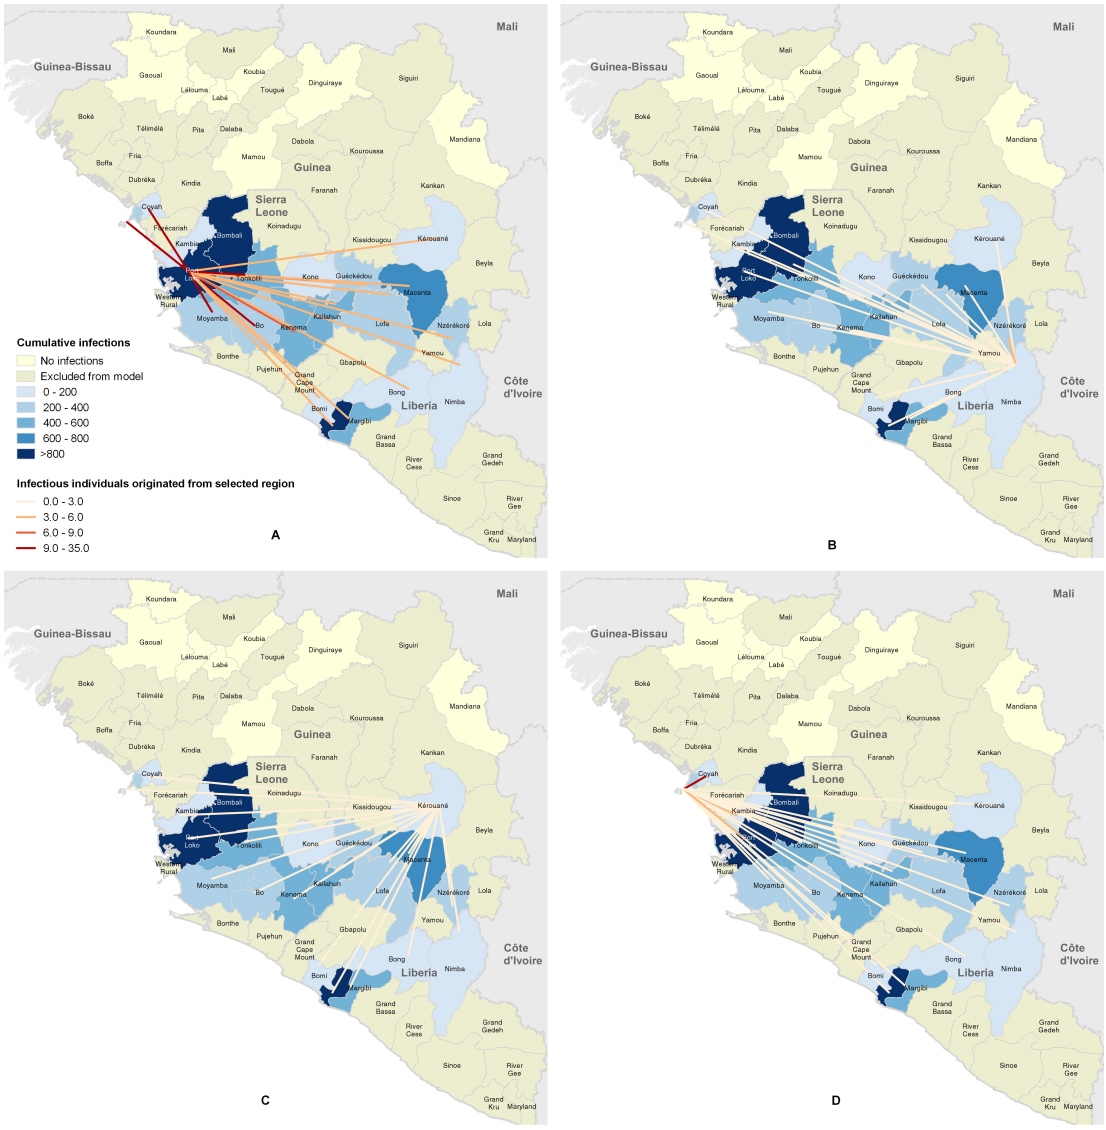

Figure 4.3: Cumulative Ebola infections from 2 September 2014 to 19 January 2015 and number of cases originating from (A) Port Loko, Sierra Leone, (B) Nimba, Liberia, (C) Kérouané, Guinea, and (D) Conakry, Guinea.

The proximity of infectious individuals to susceptible populations further contributes to the epidemic's heterogeneous spread, where some regions witness rapid increases in new infections. One of the regions closest to the epicenters of the 2014 Ebola epidemic is Kambia, Sierra Leone, where an estimated 58 infections (32% of the total in this region) originate in severely affected adjacent regions. In the coastal region of Conakry, Guinea, 68 cases (only 18% of the total) originate outside the region, due to its geographic isolation from areas with high Ebola case counts (Figure 4.3). If our model completely ignores distances between regions, the MAPE increases from 11% to 68% (using 100 days of data for calibration), demonstrating the importance of mixing patterns among susceptible and infected populations.

Most infections in a region originate from infectious individuals within the same region. However, if regions are geographically close (e.g., Port Loko and Kambia, Sierra Leone) or if they vary in disease prevalence (e.g., Bombali, Sierra Leone versus Conakry, Guinea), then the proximity between regions should be considered when allocating intervention resources. Moreover, understanding which regions are most vulnerable to Ebola transmission from neighboring regions can help policy makers anticipate where next to target relief efforts.

4.5.2 ADP Implementation

We regress different sets of state features on a sample dataset of simulated epidemic trajectories and choose the basis function with the most explanatory power: $R_{0i}S_iI_i + B, \forall i \in 1, \ldots, K$ (Table 4.2). The ADP algorithm is run for 2,000 iterations, resulting in convergence in all tested cases. The first 100 iterations are used for exploration to ensure adequate fitting of the basis function parameters; thereafter, a decreasing exploration rate of $1/0.1n$ is used. Convergence is slower for larger budgets B or higher deployment rates r as there are more potential allocation decisions to consider. Average runtime is 11 hours using one core of a 2.3-GHz Intel Xeon E5 processor.

4.5.3 Greedy R_0 versus ADP Algorithm

To illustrate the relative utility of each approach, we first estimate model parameters using periods ranging from 20 to 100 days. Assuming beds are allocated according

to *Greedy* R_0 or *ADP algorithm*, we calculate projected future Ebola cases—and the number that could be theoretically averted—compared to the forecasted epidemic with no beds available. A short fitting period of 20 or 40 days might reflect a typical scenario faced by policy makers who must quickly allocate resources during the early days of an emerging outbreak, whereas a longer fitting period represents later stages when more accurate case data are available. In our numerical experiments, all fitting periods start on September 2, 2014; resources are then allocated for the following 18 weeks, in order to fully utilize the entire budget and allow the epidemic to subside. We further test the performance of the *ADP algorithm* against the simpler *Greedy* R_0 approach for budgets between 50 and 200 beds, assuming a maximum deployment rate of five beds per week and region.

Variations in the number of days used to fit the epidemic model, or the overall budget size, highlight the relative strengths and weaknesses of each approach (Table 4.4). *Greedy* R_0 generally performs as well as or better than the *ADP algorithm* for shorter fitting periods (≤ 40 days), when the full epidemic model is less precise (MAPE $> 27\%$, see Table 4.3). *Greedy* R_0 outperforms the *ADP algorithm* by up to 44%, avoiding 1,989 infections with a budget of 50 beds and a fitting period of 40 days, compared to 1,378 avoided by *ADP algorithm*. In this case, *Greedy* R_0 allocates beds to only 5 regions (where $1.01 \leq R_{0j} \leq 1.35$), compared to 17 regions with *ADP algorithm*.

The *ADP algorithm* performs best when the underlying epidemic model is fit to a strong time series of case data: with a fitting period of 60 days and budget of 150 beds, the *ADP algorithm* averts 4,161 infections, compared to 3,876 cases averted using *Greedy* R_0, a 7% improvement. As shown in Figure 4.4, the *ADP algorithm* spreads the available budget across more regions, allocating between 2 and 23 beds in all 21 regions, as it finds a solution based on regional epidemic forecasts. In contrast, *Greedy* R_0 uniformly allocates beds to only 6 regions (where $1.06 \leq R_{0j} \leq 1.57$). With larger budgets ($> 50$ beds), the *ADP algorithm* dominates across most fitting periods; with a fitting period of 40 days *Greedy* R_0 outperforms the *ADP algorithm*, although differences are small relative to the total number of cases. This discontinuity is likely due to jumps in model parameters in regions where case numbers are not changing continuously (i.e., due to under- or over-reporting). The *ADP algorithm* considers all regions in its allocation, whereas *Greedy* R_0 considers only regions with $R_{0j} > 1$ and is therefore not sensitive to parameter changes for regions below this threshold.

Table 4.4: Number of new cases computed with the R_0 heuristic and ADP algorithm for different bed budgets and number of days of case data used for fitting the epidemic model (gray cells indicate dominant ADP solutions).

# of days used for fitting	*Greedy* R_0 Bed budget				*ADP algorithm* Bed budget			
	50	100	150	200	50	100	150	200
20	551	468	442	442	679	454	420	419
40	843	618	602	602	1455	715	662	651
60	2161	1488	1221	1184	1673	1164	936	892
80	2895	1653	1429	1348	1794	1243	970	958
100	2573	1501	1313	1270	2105	1286	999	962

4.5.4 Other Allocation Heuristics

We test the performance of the aforementioned approaches against two heuristics: *planned allocation*, based on the fraction of available beds (UNMEER, 2015) allocated to each region by participating organizations (e.g., the US Centers for Disease Control and Prevention, French Red Cross, Doctors Without Borders, Save The Children, local health ministries), and an allocation in proportion to each region's *cumulative number of infections*.

The *planned allocation* heuristic performs poorly because it fails to incorporate information about the epidemic's future trajectory. Both the *ADP algorithm* and *Greedy* R_0 outperform it, as well as the heuristic based on the *cumulative number of infections*, in every scenario tested, although the relative improvement increases with the length of the fitting period. For example, the *ADP algorithm* could prevent 3,424 infections with a budget of 50 beds and a fitting period of 60 days, while the *planned allocation* or *cumulative number of infections* heuristic could only prevent 1,424 or 2,524 infections, respectively. We find that the *ADP algorithm* offers the greatest benefit compared to the simpler heuristics when trade-offs between regions must be made. For example, with 150 available beds and 60-day fitting period, the number of infectious individuals in Conakry, Guinea exceeds the number of beds allocated; *ADP algorithm* balances the projected infections occurring in Conakry against infections in other regions if beds were assigned elsewhere.

Compared with the actual planned allocation, the *ADP algorithm* generally shifts

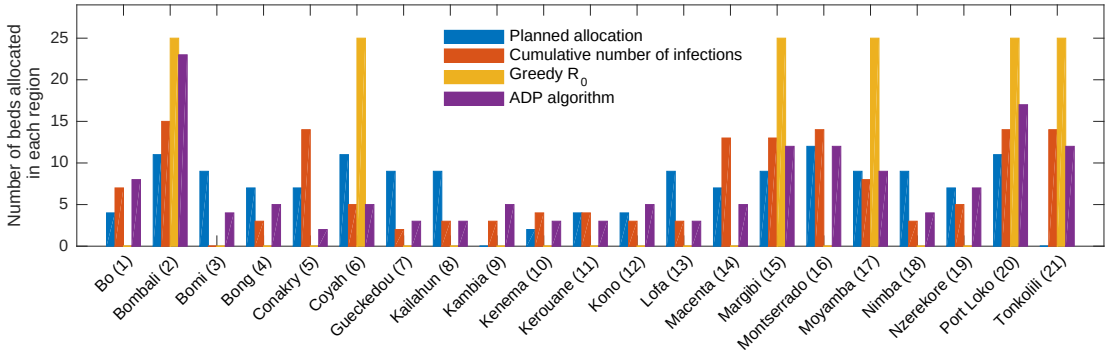

Figure 4.4: Number of beds allocated to each region (between 21 November 2014 and 27 March 2015, Budget of 150 beds, fitting period 60 days) under the approximate dynamic programming (ADP) algorithm (purple), the Greedy R_0 approach (yellow), the planned allocation (blue), and a benchmark heuristics with allocations based number of infections (red).

resources to less affected regions; this resource transfer serves to reduce the epidemic's spread from nearby, severely affected regions (Figure 4.4). For instance, the ADP solution allocates 23 ETU beds to Bombali, Sierra Leone (region 2)—more than double the planned allocation of 11 beds—because of Bombali's proximity to two highly affected regions (Tonkolili and Moyamba, Sierra Leone). However, this ADP solution comes at the cost of allocating fewer ETU beds to other regions, including Conakry, Guinea (region 5), that are farther away from highly affected regions. In contrast, the infection-based heuristic mitigates the Ebola outbreak at current "hot spots" but does not account for the nearby populations most prone to acquiring Ebola in the future, while the *Greedy* R_0 approach uniformly allocates beds across regions with a high epidemic potential, but not necessarily the greatest number of infectious individuals.

The significance of geospatial spread becomes even more evident if the epidemic model ignores disease transmission across regions (i.e., if we set $c_0 = 1$). In this case, the *ADP algorithm* results in 7.4% more infections than the scenario with inter-regional connectivity based on population proximity (i.e., if $c_0 < 1$). This finding underscores the importance of incorporating network effects into epidemic model projections and resource allocation.

Although the *ADP algorithm* and *Greedy* R_0 approach outperform other heuristics in our experiments, the quality of their solution depend on the quality of the epidemic model and associated data. Poor-quality projections based on limited data, as observed with the initial 2014 Ebola estimates for West Africa (Butler, 2014b; Meltzer

et al., 2014), are likely to result in a suboptimal allocation. Using an epidemic model that is fitted to only 10 days of data, the *ADP algorithm* would allocate 8 beds to Margibi, Liberia, leaving many other regions with significant bed shortages. Later data reveal this region to be ultimately less affected, and early projections vastly overestimate the number of cases there.

4.6 Discussion

This paper presents a novel, two-stage model to improve the allocation of limited resources across geographic regions to mitigate an emerging infectious disease outbreak. The first stage includes an epidemic model forecasting an outbreak's trajectory on a regional level. The second stage provides methods to effectively allocate intervention resources across multiple regions. We analytically derive a solution for determining the basic reproduction number, R_0, for our multi-population epidemic model. This allows us to determine each region's R_{0j}, which is subsequently used to prioritize regions where the epidemic will continue to spread (i.e., $R_{0j} > 1$) through an *Greedy* R_0 allocation scheme. We also develop a second allocation approach, the first model to combine an approximate dynamic programming algorithm with an underlying nonlinear epidemic model for computing an effective allocation of intervention resources.

We demonstrate the utility of our *ADP algorithm* and *Greedy* R_0 heuristic by first calibrating the epidemic model to data from the 2014 Ebola epidemic in West Africa and then comparing the projected number of future Ebola cases under both allocations of Ebola treatment unit beds with the actual planned allocation and with two benchmark heuristics. We find that both the *ADP algorithm* and *Greedy* R_0 heuristic identify allocations more effective than the one actually planned. The *ADP algorithm* could have prevented up to 3,424 new Ebola infections over 18 weeks compared to 1,424 prevented by the planned allocation during this period, a 58% improvement.

The epidemic model developed here features two elements that enhance its validity. First, a *connectivity matrix* of the 21 affected regions, which adjusts transmission rates between regions based on geographic distance, leads to better forecasts of a widely heterogeneous epidemic like Ebola than do extant techniques. Second, a dynamic *behavioral change parameter* for each region captures nonlinear trends in both contact patterns and self-isolation practices as the epidemic spreads differentially among com-

munities. Unlike earlier models that overestimated future Ebola cases by two orders of magnitude, our model more closely matches actual case counts—even when fewer than 100 days of data are used to estimate model parameters.

Numerical results demonstrate that *Greedy* R_0 and *ADP algorithm* are complementary methods for allocating intervention resources during emerging epidemics. While the *Greedy* R_0 approach performs well during the early days of an outbreak, the *ADP algorithm* finds better solutions later, when sufficient case data have been collected. The *Greedy* R_0 approach could therefore be relevant for decision makers in situations where case data are sparse and a quick decision needed. The *ADP algorithm* could be used when intervention resources supply has ramped up, enabling more strategic allocation decisions which consider an allocation's impact on the trajectory of the epidemic across the affected regions. Both methods incorporate improved forecast estimates of the epidemic's trajectory and can be adapted to other populations and geographic regions.

The specific compartments of the epidemic model (e.g., susceptible, exposed, vaccinated, pregnant, etc.) and associated parameters could be easily adjusted for other infectious diseases. By including other relevant interventions into the model's dynamics, our methodology could be used to identify the optimal *combination* of interventions; examples include vaccines, personal protective equipment, and insect control for mosquito-borne diseases such as Zika. The recent outbreak in Latin America has shown that modeling efforts, like the one carried out in this paper, are more than a one-time necessity. The 2014 Ebola outbreak afforded many valuable lessons, one of which was the importance of recognizing self-protective behavior as the outbreak unfolded, which ultimately reduced transmission rates and, fortunately, resulted in fewer infections than initially forecast (Heesterbeek et al., 2015). With the current Zika epidemic, behavior adaptation (e.g., through reduced pregnancy rates or human travel patterns to Zika-afflicted regions), spatial heterogeneity (e.g., due to rainfall or other climate conditions) and the effectiveness of ongoing efforts to eradicate the particular mosquito vector, all contribute to reliable forecasts of new cases (Ferguson et al., 2016). For the application to Zika, a compartmental modeling approach could be used, although the structure and interdependencies of the compartments would need to be changed to reflect the disease transmission through vectors and not humans itself. As with any infectious disease outbreak, a key challenge to reduce the disease burden of Zika is to deploy intervention efforts as fast as possible, ideally

supported by models capable of processing sparse case data.

Our study has limitations. A basic SIR model simplifies the complex natural history of disease progression; more granular health states might improve forecasts of new cases or Ebola-related deaths. We consider only a single resource in the optimal allocation; in reality, decision makers may need to allocate money across multiple resources (e.g., prevention programs, improved screening, treatment). Due to limited data, we assume a static deployment curve; one extension of our model could include dynamic deployment rate thresholds to reflect the gradual availability of resources, as when staff are trained by previously trained workers. We model connectivity among regions as proportional to the distance between capitals, which ignores the distribution of where individuals actually reside. More detailed travel behavior could illuminate contact patterns. Finally, a key challenge in reliably forecasting an emerging outbreak is the lack of high-quality data early in the epidemic. An area that merits future research is the development of a model that uses a Bayesian exploration method to accommodate stochastically arriving data points.

4.6.1 Managerial Implications

In any real-world epidemic, accurate and up-to-date data on new cases and mortality rates are limited. Our model improves the ability of policy makers to respond effectively to an emerging outbreak, providing them with more accurate forecasts and an optimized allocation of resources. One important finding from our numerical experiments is the importance of making spatially coordinated decisions. Connectivity between regions clearly affects the epidemic's spatial spread, yet these relationships are typically neglected when allocation decisions are based on such local indicators as the current number of infectious individuals in each region. Improved data sharing and coordination of epidemic relief efforts among organizations operating in different regions could help ensure that resources are prioritized to the most vulnerable populations.

When deciding how best to distribute limited resources for epidemic control, public health authorities must often balance speed of deployment with reach of the program. The availability of beds is constrained by limited financial resources and technological barriers for acquiring the necessary infrastructure and equipment, whereas a more

rapid deployment rate relies on a continuous supply of personal protective equipment and on the availability of health care workers (Butler, 2014a). Specialized training approaches, such as "train the trainer", may be necessary for ensuring that a sufficient number of adequately skilled personnel are available. During training courses, new health care workers are taught not only how to handle personal protective equipment but also how to pass their knowledge on to other new personnel. Our model results show, for example, that an increase in the deployment rate from two to eight beds per week and region would result in 47% fewer infections (with a constant budget of 150 beds and a fitting period of 60 days). Since deployment speed is crucial, it follows that preparedness for future outbreaks could be improved by adopting common standards for equipment and for training of health care workers. Furthermore, coordination and collaboration among different organizations—facilitated by institutions like the UN Office for the Coordination of Human Affairs—would almost certainly enable faster roll out of interventions to future crisis regions (Woolhouse et al., 2015).

4.6.2 Conclusion

The eradication of any infectious disease requires substantial commitments of financial and human resources. More than $1.6 billion of emergency funding was provided by different nations and institutions for the 2014 West African epidemic, yet more than 28,000 people contracted Ebola and nearly 11,000 people died of the disease. Although transmission of the Ebola outbreak has fortunately abated, there are many lessons to be learned from this tragedy. One vital lesson is that optimizing the allocation of available resources can significantly increase the number of lives saved during future infectious disease outbreaks.

Chapter 5

Summary and future research opportunities

5.1 Summary

In this thesis, we work on topics which are driven by three trends in the health care sector: (i) In countries with developed health care systems, providers and manufacturers face cost pressure as payors try to contain rising cost induced by a higher life expectancy and increasingly advanced medical procedures. (ii) With genetic diagnostic tests now being able to give insights e.g. into an individual's genetically determined cancer risk, questions arise what consequences test result should have when prophylactic treatment options are available. (iii) In the last decade, both severity and frequency of infectious disease epidemics with a potential for a global pandemic have increased, challenging current practices how government and non-government organizations perform and coordinate humanitarian interventions in affected regions. This thesis tries to provide approaches which could help tackling the above mentioned challenges through the usage of methods developed in the field of operations research. We therefore apply recent datasets to real-life problems to ensure the relevance and practicability of the approaches.

During our research on the benefits of coordination in pharmaceutical supply chains, we examine if improved coordination between partners in a pharmaceutical supply chain could help drive down supply chain cost, a significant cost driver in this area, while maintaining the overall service level. We use demand data from the German pharmaceutical market and model a supply chain using a discrete-event simulation model fed with parameters from German pharmacies and wholesalers. To analyze the effect of coordination on supply chain cost and service level, we construct three scenar-

ios which we later compare through using MANOVA: (1) No coordination, (2) sharing of point-of-sale demand information, and (3) full coordination, including joint decision of order quantities in all echelons, also known as collaborative planning, forecasting and replenishment (CPFR). Each of the scenarios is tested by varying internal and external circumstances like varying demand patterns, product shelf life, production capacity and inventory allocation scheme. Results show that CPFR dominates the other two scenarios with the lowest cost and highest fill rate under all circumstances tested, although its benefit depends on the characteristics of the SC. Benefits of full coordination are highest when demand variation is high and product shelf life is low, suggesting an implementation for seasonal and low shelf-life products. A disruption common for pharmaceutical supply chains, constrained production capacity, is not mitigated better through more intensive coordination. Further results suggest that although CPFR is the dominant coordination method tested, an implementation of point-of-sale information sharing only might be more beneficial as it already captures on average 89% of CPFR's benefit while incurring less implementation efforts.

In our work on the optimization of cancer prevention strategies for BRCA carriers, we provide insights in how women carrying a mutation could maximize their QALYs through an optimal timing of prophylactic surgeries. The model is calibrated with data from recent clinical studies and could help patients and health care practitioners to make trade-offs between different treatment options which reduce the elevated risk of breast- and ovarian cancer. We analyze if and when performing a surgery would maximize the accumulated and discounted QALYs of a BRCA1/2 mutation carrier. Results with a baseline parameter set show that a BRCA1 carrier should perform a BM between age 30 and 60. Before and afterwards, the loss in quality-of-life through the surgery is not offset by the reduction in breast cancer risk. Additionally, a BSO is QALY-maximizing from the age of 40. Even though the benefit of a BSO reduces with a carrier's age, its impact on the quality of life after 50 is low compared to the BM as one important factor, its negative impact on the ability to reproduce, is not relevant after the age of 50 in most cases. Following the QALY-maximizing surgery sequence, simulation results of the model indicate that a BRCA1 carrier would increase her average life expectancy from 70 to 77 years. For BRCA2 carriers, QALY-maximizing surgery ages are pushed back as the overall breast and ovarian caner risk is lower than for BRCA1 carriers. Sensitivity analysis show that the recommended surgery age is heavily dependent on a carriers preferences about the impact of the surgeries on her quality-of-life. Besides the optimal numerical solution, we provide insights into the

structural properties of a simplified version of our model. Assuming a post-cancer scenario, where either the breasts or ovaries have already been removed as part of the cancer treatment, an optimal control threshold for cancer risk exists, after which it is always optimal to remove the remaining organ.

Two methods for the allocation of intervention resources in regions affected by infectious disease outbreaks based on an epidemic model are presented in our work with data from the 2014 Ebola outbreak in West Africa. We first calibrate a multi-population epidemic model with a novel behavioral dampening coefficient to data from the 2014 Ebola outbreak in West Africa. We choose ETU beds as the intervention resource to be allocated across 21 regions in the three countries affected by the Ebola outbreak: Guinea, Liberia, and Sierra Leone. Our first allocation approach, Greedy R_0, is based on prioritizing the regions by their basic reproduction, R_0, which we determine based on a developed analytical solution. The second allocation approach uses ADP to compute an efficient allocation of beds across the affected regions, as the state space is too large to compute an optimal solution through DP. When comparing both approaches with one other heuristics and the actual allocation, results show that neither the Greedy R_0 nor the ADP algorithm provide a dominant solution as for small bed budgets and imprecise epidemic models, the Greedy R_0 approach provides the most effective allocation. For more complex allocation decisions where larger bed budgets are involved and precise epidemic model parameters are available, the ADP algorithm proves to be the dominant solution as it incorporates the impact of the allocated beds on the future trajectory of the epidemic in its decision. By applying these two approaches, our model results suggest that up to 2,000 infections could have been saved compared with the actually planned allocation.

5.2 Future research

With the results obtained in each of the three research projects included in this thesis, we enable further research in numerous areas. The following paragraphs elaborate on possible extensions considering future trends in each of the fields covered.

With the increase in hardware and software capabilities of analyzing large amounts of data, the technical barriers for exchanging data and coordinate planning in a SC get lower. Implementation of a coordinated planning scheme is remaining complicated

though, with legal and behavioral boundaries remaining strong. Further research efforts therefore could explore e.g. how incentive systems for SC partners have to be designed to prevent inefficient behavior of single entities. While lack of mutual trust in a joint planning system will almost always be a barrier preventing the full integration of a SC, future research could also elaborate on the effects of a participation of only a share of all SC partners in a collaboration scheme.

Due to sinking cost for genetic testing and increased awareness, the number of precancer BRCA1/2 detections is likely to grow. This does not only require intensified efforts to develop patient-specific recommendations about the options to be considered after a detection, but will also magnify the number of clinical studies providing insights BRCA-specific parameters which could be used in a decision support model. An important area of future research will therefore be the further development of decision support model increasing both the level of detail and the quality of the input parameters. Another important area is the dissemination of the model result beyond the academic literature. As most patients do not have the necessary skills and knowledge to understand the results presented in an academic research paper, recommendations need to be transformed in a form understandable by average patients without sacrificing the clarity about the model's inherent uncertainty.

During infectious disease outbreaks, data is available while the epidemic is still growing and allocation decisions are made. Further developments in the enhancement of epidemic models and resource allocation recommendations should therefore concentrate on using the available data to compute simple, but actionable methods for adapting existing models to emerging situations. This could e.g. be achieved through a modular approach in epidemic modeling, where a multi-population module is only used when an outbreak shows signs of spatial heterogeneity. Further work could also advance the quality of solutions computed through ADP. This could include the application of additional exploration methods as well as the integration of additional basis function. To further increase the impact in this field of research, future efforts should also consider the standards and procedures applied by organizations deploying the allocated resources.

Appendix A

Appendix to Chapter 2

A.1 Demand dataset

The available demand data contains demand time series for pharmaceuticals from five different markets (Allergy, Cold, Diabetes, Hypertension, and Parkinson disease) and the corresponding specification of the pharmaceuticals (e.g. wholesaler price, shelf life, etc.). Demand data was collected and provided by *Insight Health*, pharmaceuticals were matched with specification provided by *ABDATA* through their PNZ (pharmaceutical identification number). Pharmaceuticals were included in the simulation and selected through the following steps:

1) All pharmaceuticals without any demand were dropped

2) All pharmaceuticals with a cumulative demand of 1,000 units or less were dropped

3) All pharmaceuticals with data missing for one or more state(s) were dropped

4) The pharmaceuticals were sorted by their variance, the products with the highest, lowest, and median variance were selected.

Figure A.1 shows the demand time series for the pharmaceuticals with stable and seasonal demand. To make the simulation computationally feasibly, only 3 of the 16 available states were used. Hamburg, Rhineland-Palatinate, and Mecklenburg-Western Pomerania, represent about 9% of the total German population.

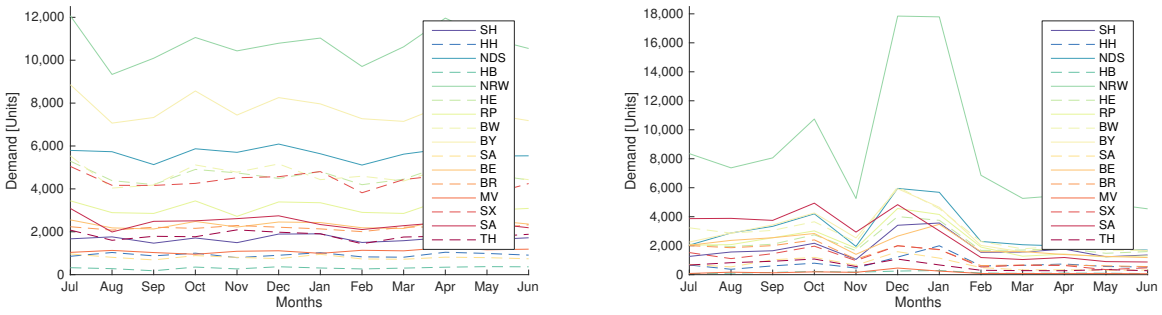

Figure A.1: Demand pattern of pharmaceutical with stable (left) and seasonal (right) demand

A.2 Detailed MANOVA results

A.2.1 Estimated DV means

Table A.1 - A.6 contain the mean of the dependent variables, estimated after the significance of the differences between the treatment groups were confirmed through the MANOVA.

Table A.1: Estimated SCC and FR means using different coordination schemes (H1)

Dependent variables	SCS	Mean	Boundaries (95 % CI)	
			Lower	Upper
SCC	CONV	408.1956	405.994	410.3972
	POS	265.7893	263.649	267.9301
	CPFR	248.0453	245.9296	250.1609
FR	CONV	0.6742	0.6721	0.6764
	POS	0.8661	0.8640	0.8682
	CPFR	0.8967	0.8947	0.8987

Table A.2: Estimated SCC and FR means using different coordination schemes and demand variabilities (H2)

Dependent variables	PRO	SCS	Mean	Boundaries (95 % CI)	
				Lower	Upper
SCC	1	CONV	384.2577	380.6108	387.9045
		POS	237.7617	233.9320	241.5914
		CPFR	242.9663	239.3218	246.6109
	2	CONV	406.4219	402.5152	410.3286
		POS	288.1984	284.5538	291.8429
		CPFR	254.6082	250.9590	258.2573
	3	CONV	434.3498	430.4666	438.2330
		POS	271.8097	268.1651	275.4543
		CPFR	246.6222	242.9272	250.3172
FR	1	CONV	0.7366	0.7331	0.7402
		POS	0.9592	0.9555	0.9629
		CPFR	0.9637	0.9602	0.9672
	2	CONV	0.6000	0.5962	0.6038
		POS	0.8400	0.8365	0.8436
		CPFR	0.9174	0.9139	0.9209
	3	CONV	0.6853	0.6816	0.6891
		POS	0.7975	0.7940	0.8010
		CPFR	0.8077	0.8041	0.8112

Table A.3: Estimated SCC and FR means using different coordination schemes and product shelf lifes (H3).

Dependent variables	SL	SCS	Mean	Boundaries (95 % CI)	
				Lower	Upper
				Lower	Upper
SCC	12	CONV	332.3801	328.4202	336.3401
		POS	258.0682	254.4055	261.7310
		CPFR	243.5934	239.9350	247.2517
	6	CONV	355.5414	351.8901	359.1926
		POS	264.1034	260.3693	267.8375
		CPFR	243.7415	240.0663	247.4167
	3	CONV	536.8093	532.9823	540.6364
		POS	275.1304	271.4206	278.8403
		CPFR	256.8359	253.1822	260.4895
FR	12	CONV	0.6261	0.6223	0.6299
		POS	0.8687	0.8652	0.8723
		CPFR	0.8899	0.8863	0.8934
	6	CONV	0.6971	0.6935	0.7006
		POS	0.8705	0.8669	0.8741
		CPFR	0.9006	0.8970	0.9041
	3	CONV	0.6984	0.6948	0.7021
		POS	0.8591	0.8555	0.8626
		CPFR	0.8995	0.8960	0.9030

Table A.4: Estimated SCC and FR means using different coordination schemes and production capacity (H4).

Dependent variables	Capacity	SCS	Mean	Boundaries (95 % CI)	
				Lower	Upper
SCC	0.9 × Demand	CONV	362.8908	359.1873	366.5942
		POS	246.9107	243.2631	250.5582
		CPFR	227.7218	224.0248	231.4189
	1.1 × Demand	CONV	422.0921	418.2689	425.9153
		POS	270.8907	267.1469	274.6345
		CPFR	255.2361	251.5909	258.8814
	1.3 × Demand	CONV	441.2523	437.3498	445.1548
		POS	279.9230	276.1957	283.6503
		CPFR	261.8260	258.1807	265.4713
FR	0.9 × Demand	CONV	0.6505	0.6469	0.6541
		POS	0.8044	0.8009	0.8079
		CPFR	0.8202	0.8166	0.8237
	1.1 × Demand	CONV	0.6717	0.6680	0.6754
		POS	0.8860	0.8824	0.8897
		CPFR	0.9244	0.9208	0.9279
	1.3 × Demand	CONV	0.7014	0.6976	0.7052
		POS	0.9120	0.9084	0.9157
		CPFR	0.9481	0.9445	0.9516

Table A.5: Estimated SCC and FR means using different coordination and allocation schemes (H5).

Dependent variable	Allocation Policy	SCS	Mean	Boundaries (95 % CI)	
				Lower	Upper
FR	FIFO	CONV	0.6674	0.6636	0.6711
		POS	0.8928	0.8892	0.8963
		CPFR	0.8984	0.8949	0.9019
	Proportional	CONV	0.6972	0.6935	0.7008
		POS	0.8776	0.8739	0.8813
		CPFR	0.9076	0.9040	0.9111
	Practical	CONV	0.6587	0.6551	0.6624
		POS	0.8286	0.8250	0.8321
		CPFR	0.8844	0.8809	0.8880

Table A.6: Estimated manufacturer and pharmacy cost shares using different coordination schemes (H6)

Dependent variables	SCS	Mean	Boundaries (95 % CI)	
			Lower	Upper
CSM	CONV	0.4732	0.470358	0.4760
	POS	0.7761	0.773	0.7789
	CPFR	0.8455	0.842778	0.8482
CSPH	CONV	0.0378	0.0375	0.0380
	POS	0.0742	0.0740	0.0745
	CPFR	0.0798	0.0795	0.0801

A.2.2 ANOVA results

In Table A.7 - A.10, the results of the five ANOVAs are presented as follow-up analysis to the MANOVA assumptions discussed in Section 2.5. Results show that also for each univariate ANOVA, the relevant effects are significant at a .01 significance level.

Table A.7: ANOVA results with SCC as dependent variable.

Effect	Sum of squares	Mean square	F value	Prob > F
SCS	35070020.18	17535010.09	6262.18	0.0000
ALL	411583.55	205791.78	73.49	0.0000
PRO × SCM	2017756.15	336292.69	120.10	0.0000
CAP × SCM	3480677.00	580112.83	207.17	0.0000
SL × SCM	18568242.01	3094707.00	1105.20	0.0000
ALL × SCM	614436.29	153609.07	54.86	0.0000

Table A.8: ANOVA results with FR as dependent variable.

Effect	Sum of squares	Mean square	F value	Prob > F
SCS	66.36	33.18	12738.12	0.0000
ALL	1.76	0.88	338.02	0.0000
PRO × SCM	28.02	4.67	1792.81	0.0000
CAP × SCM	13.19	2.20	844.23	0.0000
SL × SCM	2.55	0.43	163.41	0.0000
ALL × SCM	0.85	0.21	81.14	0.0000

Table A.9: ANOVA results with CSPH as dependent variable.

Effect	Sum of squares	Mean square	F value	Prob > F
SCM	2.37	1.19	23431.27	0.0000
ALL	0.04	0.02	346.29	0.0000
PRO × SCM	0.25	0.04	807.75	0.0000
CAP × SCM	0.00	0.00	7.26	0.0000
SL × SCM	0.01	0.00	33.36	0.0000
ALL × SCM	0.03	0.01	154.81	0.0000

Table A.10: ANOVA results with CSM as dependent variable.

Effect	Sum of squares	Mean square	F value	Prob > F
SCS	178.64	89.32	19348.73	0.0000
ALL	2.00	1.00	216.96	0.0000
PRO × SCM	13.13	2.19	474.13	0.0000
CAP × SCM	0.42	0.07	15.24	0.0000
SL × SCM	5.86	0.98	211.42	0.0000
ALL × SCM	1.72	0.43	93.27	0.0000

Appendix B

Appendix to Chapter 3

B.1 Structural properties

In this Section, we first provide the remaining Propositions and Lemmas which have been omitted in Section 3.4.1 and give the necessary proofs afterwards.

B.1.1 Supporting Propositions and Lemmas

In Proposition B.1, we show that $v_t(s_t)$ is nonincreasing in s_t, implying that a carrier's expected QALYs do not increase with her age and cancer risk, respectively.

Proposition B.1. *If the transition probability matrix for action W is IFR for all $t = 1, 2, \ldots, T$, then $v_t(s_t)$ is nonincreasing in s_t, for $s_t = x^{20}, x^{21}, \ldots, x^{AGE}$, and $t = 1, 2, \ldots, T-1$.*

To show that $v_t(s_t)$ is also nonincreasing in t, which means that the expected QALYs of a patient do not increase with her age, we first present Lemma B.1:

Lemma B.1. *If Assumption 3.7 holds for $t = 1, 2, \ldots, T-1$, then for any $f(i)$ nonincreasing in i, the following holds:*

$$\sum_{s'_t \in \hat{S}} p_t(s'_t \mid i) f(s'_t) \geq \sum_{s'_t \in \hat{S}} p_{t+1}(s'_t \mid i) f(s'_t). \tag{B.1}$$

Proposition B.2. *The optimal value function, $v_t(s_t)$, is nonincreasing in t for all $s_t \in \hat{S}$.*

To prove Theorem 3.1, the following Lemma is used:

Lemma B.2. *$\mathbb{P} = [p_t(j \mid i)]$ is an IFR transition probability matrix for $i, j = 1, 2, \ldots, N$, such that $\sum_{k=i+1}^{k^*} p_t(k \mid i+1) \geq \sum_{k=i+1}^{k^*} p_t(k \mid i)$ for $i < k^* \leq N$ and $t = 1, 2, \ldots, T-1$. If $f(i)$ is a nonincreasing function in i, then the following holds:*

$$\sum_{k=1}^{i} \{p_t(k \mid i) - p_t(k \mid i+1)\} f(k) \geq \sum_{k=1}^{i} \{p_t(k \mid i) - p_t(k \mid i+1)\} f(i), \tag{B.2}$$

$$\sum_{k'=i+1}^{k^*} \{p_t(k' \mid i) - p_t(k' \mid i+1)\} f(k') \geq \sum_{k'=i+1}^{k^*} \{p_t(k' \mid i) - p(k' \mid i+1)\} f(i+1). \tag{B.3}$$

B.1.2 Proofs

Proof of Proposition B.1 The proof of Proposition B.1 is similar to the proof of Proposition 4.7.3 in Puterman (2014) and therefore omitted here.

Proof of Lemma B.1 The sum over a row in the transition probability matrix results to $\sum_{s'=0}^{x^{death}} p_t(s' \mid i) = \sum_{s'=0}^{x^{death}} p_{t+1}(s' \mid i) = 1$. Let

$$\begin{aligned}
E &= \sum_{s'=x^{20}}^{x^{death}} p_{t+1}(s' \mid i) - \sum_{s'=x^{20}}^{x^{death}} p_t(s' \mid i) = 0, \text{ then} \\
E &= \left\{ \sum_{s'=x^{20}}^{x^{death}} p_{t+1}(s' \mid i) - \sum_{s'=x^{20}}^{x^{death}} p_t(s' \mid i) \right\} f(x^{20}) \\
&\geq p_{t+1}(x^{20} \mid i) f(x^{20}) - p_t(x^{20} \mid i) f(x^{20}) \\
&\quad + \left\{ \sum_{s'=x^{21}}^{x^{death}} p_{t+1}(s' \mid i) - \sum_{s'=1}^{x^{death}} p_t(s' \mid i) \right\} f(x^{21}) \qquad \text{(B.4)} \\
&\geq \sum_{s'=x^{20}}^{x^{21}} p_{t+1}(s' \mid i) f(s') - \sum_{s'=x^{20}}^{x^{21}} p_t(s' \mid i) f(s') \\
&\quad + \left\{ \sum_{s'=x^{22}}^{x^{death}} p_{t+1}(s' \mid i) - \sum_{s'=x^{22}}^{x^{death}} p_t(s' \mid i) \right\} f(x^{22}). \qquad \text{(B.5)}
\end{aligned}$$

(B.4) and (B.5) hold through the assumption that $f(i)$ is nonincreasing in i. It therefore follows that:

$$E \leq \sum_{s'=x^{20}}^{x^{death}} p_{t+1}(s' \mid i) f(s') - \sum_{s'=x^{20}}^{x^{death}} p_t(s' \mid i) f(s') \tag{B.6}$$

$$\geq \sum_{s'=x^{20}}^{x^{death}} p_{t+1}(s' \mid i) f(s') - \sum_{s'=x^{20}}^{x^{death}} p_t(s' \mid i) f(s'),$$

which leads to $\sum_{s' \in \hat{S}} p_t(s' \mid i) f(s') \geq \sum_{s' \in \hat{S}} p_{t+1}(s' \mid i) f(s')$.

Proof of Proposition B.2 We use backward induction to proof this Proposition: For $t = T - 1$:

$$\begin{aligned} v_{T-1}(s_{T-1}) &\geq r_{T-1}(s_{T-1}, BSO) \\ &\geq r_T(s_T, BSO) && \text{(B.7)} \\ &= v_T(s_T). && \text{(B.8)} \end{aligned}$$

Assumption 3.2 and 3.3 can be used to deduce (B.7) and the boundary condition of given in Equation (3.2) leads to (B.8). It follows that $v_{T-1}(s_{T-1}) \geq v_T(s_T)$, which is a sufficient proof for the base case. The Proposition is assumed to hold as well for $t = t_0 \ \forall s_t \ \hat{S}$. That the theorem holds for $t = t_0 - 1$ is proven through:

$$\begin{aligned} v_{t_0-1}(s_t) = \max \Bigg\{ & r_{t_0-1}(s_{t_0-1}, BSO) + \gamma \sum_{s' \in \hat{S}} p_{t_0-1}(s' \mid s_{t_0-1}, BSO) v_{t_0}(x'), \\ & r_{t_0-1}(s_{t_0-1}, W) + \gamma \sum_{s' \in \hat{S}} p_{t_0-1}(s' \mid s_{t_0-1}, W) v_{t_0}(s') \Bigg\} \\ \geq \max \Bigg\{ & r_{t_0}(s_{t_0}, BSO) + \gamma \sum_{s' \in \hat{S}} p_{t_0}(s' \mid s_{t_0}, BSO) v_{t_0}(s'), \\ & r_{t_0}(s_{t_0}, W) + \gamma \sum_{s' \in \hat{S}} p_{t_0}(s' \mid s_{t_0}, W) v_{t_0}(s') \Bigg\} && \text{(B.9)} \\ \geq \max \Bigg\{ & r_{t_0}(s_{t_0}, BSO) + \gamma \sum_{s' \in \hat{S}} p_{t_0}(s' \mid s_{t_0}, BSO) v_{t_0+1}(s') \\ & r_{t_0}(s_{t_0}, W) + \gamma \sum_{s' \in \hat{S}} p_{t_0}(s' \mid s_{t_0}, W) v_{t_0+1}(s') \Bigg\} && \text{(B.10)} \\ = v_{t_0}(s_{t_0}), && \text{(B.11)} \end{aligned}$$

where (B.9) follows from Assumption 3.1 - 3.3 and Lemma B.1. The induction hypothesis $v_{t_0}(s_t) \geq v_{t_0+1}(s_t)$ leads to (B.10). The Proposition therefore holds for all $t = 1, 2, \ldots, T - 1$.

Proof of Lemma B.2 The following proof of Lemma B.2 is given for the infinite case in Alagoz et al. (2004). To proof Equation B.2, we repeat that the IFR assumption implies that $\sum_{j=1}^{k} p_t(j \mid i) \geq \sum_{j=1}^{k} p_t(j \mid i+1)$ for any $k \in \hat{S}$. Let

$$
\begin{aligned}
&\sum_{k=1}^{i}\{p_t(k \mid i) - p_t(k\ i+1)\}f(k) \\
&\quad = \{p_t(x^{21} \mid i) - p_t(x^{21} \mid i+1)\}f(x^{21}) + \sum_{k=x^{22}}^{i}\{p_t(k \mid i) - p_t(k \mid i+1)\}f(k) \\
&\quad \geq \{p_t(x^{21} \mid i) - p_t(x^{21} \mid i+1)\}f(x^{22}) + \sum_{k=x^{22}}^{i}\{p_t(k \mid i) - p_t(k \mid i+1)\}f(k) \qquad \text{(B.12)} \\
&\quad = \{p_t(x^{21} \mid i) + p_t(x^{22} \mid i) - p_t(x^{21} \mid i+1) - p_t(x^{22} \mid i+1)\}f(x^{22}) \\
&\qquad + \sum_{k=x^{23}}^{i}\{p_t(k \mid i) - p_t(l \mid i+1)\}f(k) \\
&\quad \geq \{p_t(x^{21} \mid i) + p_t(x^{22} \mid i) - p_t(x^{21} \mid i+1) - p_t(x^{22} \mid i+x^{22})\}f(x^{23}) \\
&\qquad + \sum_{k=x^{23}}^{i}\{p_t(k \mid i) - p_t(k \mid i+1)\}f(k), \qquad \text{(B.13)}
\end{aligned}
$$

with (B.12) following from $p_t(x^{21} \mid i) \geq p_t(x^{21} \mid i+1)$ and $f(x^{21}) \geq f(x^{22})$. Inequality (B.13) holds as $p_t(x^{21} \mid i) + p_t(x^{22} \mid i) \geq p_t(x^{21} \mid i+1)$ and $v(x^{21}) \geq v(x^{22})$. The complete proof follows from the application of the same procedure to the remaining states $x^{23}, \ldots, x^{death}$. Equation (B.3) requires a similar proof which is omitted here.

Proof of Theorem 3.1 The following inequalities hold if $a^*(s_t)$ exists:

$$v_t(\overline{s_t}, BSO) \geq r_t(\overline{s_t}, W) + \gamma \sum_{s' \in \hat{S}} p_t(s' \mid \overline{s_t}, W) v_{t+1}(s'), \qquad \text{(B.14)}$$

$$v_t(\overline{s_t}+1, BSO) \geq r_t(\overline{s_t}+1, W) + \gamma \sum_{s' \in \hat{S}} p_t(s' \mid \overline{s_t}+1, W) v_{t+1}(s'). \qquad \text{(B.15)}$$

Assume $a^*(s_t) = BSO$ for $s_t = \overline{s_t}$ and $a^*(s_t+1) = W$ only for $s_t = \overline{s_t}+1$, and $t = 0, 1, \ldots, T-1$. It follows:

$$
\begin{aligned}
&v_t(\overline{s_t}, BSO) \geq r_t(\overline{s_t}, W) + \gamma \sum_{s' \in \hat{S}} p_t(s' \mid \overline{s_t}, W) v_{t+1}(s') \\
&v_t(\overline{s_t}, BSO) < r_t^*(\overline{s_t}+1, W) \\
&\qquad\qquad + \gamma \sum_{s' \in \hat{S}} p_t(s' \mid \overline{s_t}+1, W) v_{t+1}(s')
\end{aligned}
$$

$$
\begin{aligned}
&v_t(\overline{s_t}, BSO) - \\
&v(\overline{s_t}+1, BSO) > r_t(\overline{s_t}, W) - r_t(\overline{s_t}+1, W) \\
&\qquad + \gamma \sum_{s' \in \hat{S}} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(s') \\
&\qquad \geq \gamma \sum_{s' \in \hat{S}} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(s') \qquad \text{(B.16)} \\
&\qquad = \gamma \sum_{s'=0}^{\overline{s_t}} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(s') \\
&\qquad\quad + \gamma \sum_{s''=\overline{s_t}+1}^{x^{AGE}} \{p_t(s'' \mid \overline{s_t}, W) - p_t(s'' \mid \overline{s_t}+1, S)\} v_{t+1}(s'') \\
&\qquad \geq \gamma \sum_{s'=0}^{s_t*} \{p_t(s' \mid s_t*, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}) \\
&\qquad\quad + \gamma \sum_{s''=\overline{s_t}+1}^{x^{AGE}} \{p_t(s'' \mid \overline{s_t}, W) - p_t(s'' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}+1) \qquad \text{(B.17)} \\
&\qquad \geq \gamma \sum_{s'=0}^{s_t*} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}+1) \\
&\qquad\quad + \gamma \sum_{s''=\overline{s_t}+1}^{x^{AGE}} \{p_t(s'' \mid \overline{s_t}, W) - p_t(s'' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}+1) \qquad \text{(B.18)} \\
&\qquad = \gamma \sum_{s'=0}^{x^{AGE}} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}+1) \\
&\qquad \geq \gamma \sum_{s'=0}^{x^{AGE}} \{p_t(s' \mid \overline{s_t}, W) - p_t(s' \mid \overline{s_t}+1, W)\} v_{t+1}(\overline{s_t}+1, BSO) \qquad \text{(B.19)} \\
&\qquad = \gamma \{p_t(x^{death} \mid \overline{s_t}+1, W) \\
&\qquad\quad - p_t(x^{death} \mid \overline{s_t}, W)\} r_{t+1}(\overline{s_t}+1, BSO) \qquad \text{(B.20)}
\end{aligned}
$$

Inequality (B.16) follows from Assumption 3.1 as well as $p_t(j \mid s, W) = 0$ for $j \in \{x^{surg}\}$ and $s \in \hat{S} \backslash x^{surg}$. Inequality (B.17) follows from Proposition B.1 and Lemma B.2. Being IFR, $\mathbb{P}_t^W$ implies that $\sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t, W) \leq \sum_{s'=s_t+1}^{x^{death}} p_t(s' \mid s_t+1, W)$ and $\sum_{s'=0}^{s_t} p_t(s' \mid s_t, W) \geq \sum_{s'=s_t}^{s_t} p_t(s' \mid s_t+1, W)$. By using Proposition B.1 to state that $v_{t+1}(s_t) \geq v_{t+1}(s_t+1)$, $v_{t+1}(s_t)$ in Equation (B.17) can be replaced by $v_{t+1}(s_t+1)$ in Equation (B.18). With $\mathbb{P}_t^W$ being IFR and $v_{t+1}(s_t+1) \geq r_{t+1}(s_t+1, BSO)$, we can conclude inequality (B.19). Equation (B.20) can therefore be replaced by:

$$
\frac{v_t(\overline{s_t}, BSO) - v_t(\overline{s_t}+1, BSO)}{v_{t+1}(\overline{s_t}+1, BSO)} > \gamma \{p_t(x^{death} \mid \overline{s_t}+1, W) - p_t(x^{death} \mid \overline{s_t}, W)\}, \qquad \text{(B.21)}
$$

which contradicts Equation (3.1), from which the proof follows.

B.2 Data format conversion

This section describes the formulas used to convert the data collected from different clinical studies into a form which is usable in the MDP.

Cancer risk Chen & Parmigiani (2007) provide the probability of being diagnosed with breast- and ovarian cancer in 10-year intervals. As we use 1-year decision epochs in the MDP, these probabilities need to be converted to yearly probabilities. We do this through the following Equations, with p_{10} being the probability of an event happening during a 10-year interval, and p_1 the probability of the same event happening during for a 1-year interval:

$$p_1 = 1 - \sqrt[10]{1 - p_{10}} \tag{B.22}$$

Risk reduction given as odds ratio (OR) The reduction of cancer risk through prophylactic surgery is given as OR by Eisen et al. (2005). An odds ratio is defined in the following Equation with p_b as the baseline probability of an event and p_{rr} as the probability of this event after the risk reducing surgery:

$$OR = \frac{\frac{p_{rr}}{1-p_{rr}}}{\frac{p_b}{1-p_b}}. \tag{B.23}$$

Being in line with Grant (2014), the probability after having performed a risk reducing surgery therefore results to:

$$p_{rr} = \frac{ORp_b}{1 - p_b + ORp_b}. \tag{B.24}$$

Risk reduction given as hazard ratio (HR) Rebbeck et al. (2004) and Rebbeck et al. (2009) use HRs as indicator for the reduction in cancer risk through prophylactic surgery. We convert the probability p of being diagnosed with cancer in t years to a

rate r (assuming a constant rate) through:

$$r = -\ln(1-p)t^{-1} \tag{B.25}$$

The rate r is then adjusted for the risk reduction through surgery by multiplying it with the HR.

B.3 Model verification and validation

To ensure that the model described in Section 3.3 matches reality in the intended way, we follow model verification and validation steps proposed by Gass (1983).

B.3.1 Model verification

Model verification aims to ensure that the model "runs as intended", in our case meaning that the written software matches the mathematical expressions given in the model description. The methods used during the model development process are common in software development projects:

- **Modular coding:** We start coding a very simple MDP containing only one variable and test it using example parameters. Once accurate performance has been assured, we add the next variable and test the model behavior with an extended set of example parameters. This procedure is repeated until the MDP is completed.
- **Documentation:** To document the content of the algorithm and to make it easily interpretable for people not involved in the research project we comment the non-self-explaining lines of code while writing the algorithms.
- **Model output verification:** To ensure the correctness of the full model, we conduct extensive sensitivity analysis not only with parameter values relevant for answering the research question but also indicating when modifying the input parameters the model results behave in the intended way. We also perform simulations of the MDP to ensure that the input parameters are processed in the intended way.

B.3.2 Model validation

According to Gass (1983), model validation "tests the agreement between the behavior of the model and the real world system being modeled". We apply validation techniques defined by Gass (1983) as well as Sargent (2013) to our model:

- **Comparison to other models:** We compare the outputs of our model with existing publications in the research area (e.g. Abdollahian & Das, 2015; Grann et al., 1998; van Roosmalen et al., 2002). Although our results unsurprisingly do not exactly match the result of the aforementioned works, the general direction of the results is comparable.
- **Data validity:** While most of our parameter values have been empirically estimated during clinical studies, we use only recent studies with a sufficient population size (see also Section 3.5.1).
- **Face validity:** The input parameters and assumptions are discussed with clinical expert in the field of breast and ovarian cancer, e.g. from UCLA University hospital or the University Medical Center Hamburg-Eppendorf (UKE).
- **Extreme condition tests:** We test the model output behavior through setting the input parameters to extreme levels, e.g. assuming a QALY impact of prophylactic surgeries of 0 or 1 or setting the risk of one cancer to 0 or 1.
- **Logical / mathematical validity:** The model structure is documented for validation in Section 3.3 and analyzed in Section 3.4.1. MDP has been applied numerous times to similar research questions (see Section 3.2).
- **Sensitivity analysis:** We run a wide set of sensitivity analysis to proof the robustness of our approach under different parameter assumptions (see Section 3.5).

Appendix C

Appendix to Chapter 4

C.1 Paper 3 Appendix

We provide additional model details and results on both the ADP algorithm C.2 and the calibration of the epidemic model C.3. An overview about the relationships between the different model elements and the data used is presented in Fig. C.1.

C.2 ADP Algorithm

The pseudocode of the ADP policy iteration algorithm is presented in Fig. C.2. The mechanics of the Knapsack formulation used to efficiently compute an approximate solution of row 13 of the ADP algorithm are presented in Section C.2.1. Details on the recursive updating of the parameter vector $\boldsymbol{\theta}_t^n$, performed in line 19, are presented in section C.2.2. We provide details on the used exploration method in Section C.2.3.

C.2.1 Knapsack Formulation.

The resource allocation step in the ADP algorithm can be performed only for problems with small action spaces, since sweeps over the whole action space are computationally expensive. Although this problem is similar to a Knapsack problem, most known solution methods cannot be applied because the objective function is neither convex nor concave (Kellerer et al., 2004). We therefore divide each allocation problem into r subproblems, where r is the maximum number of additional beds which can be allocated per region per period. In each of the resulting subperiods, indexed by τ, only 0 or 1 beds can be allocated to each region. The resulting 0/1 nonlinear Knapsack

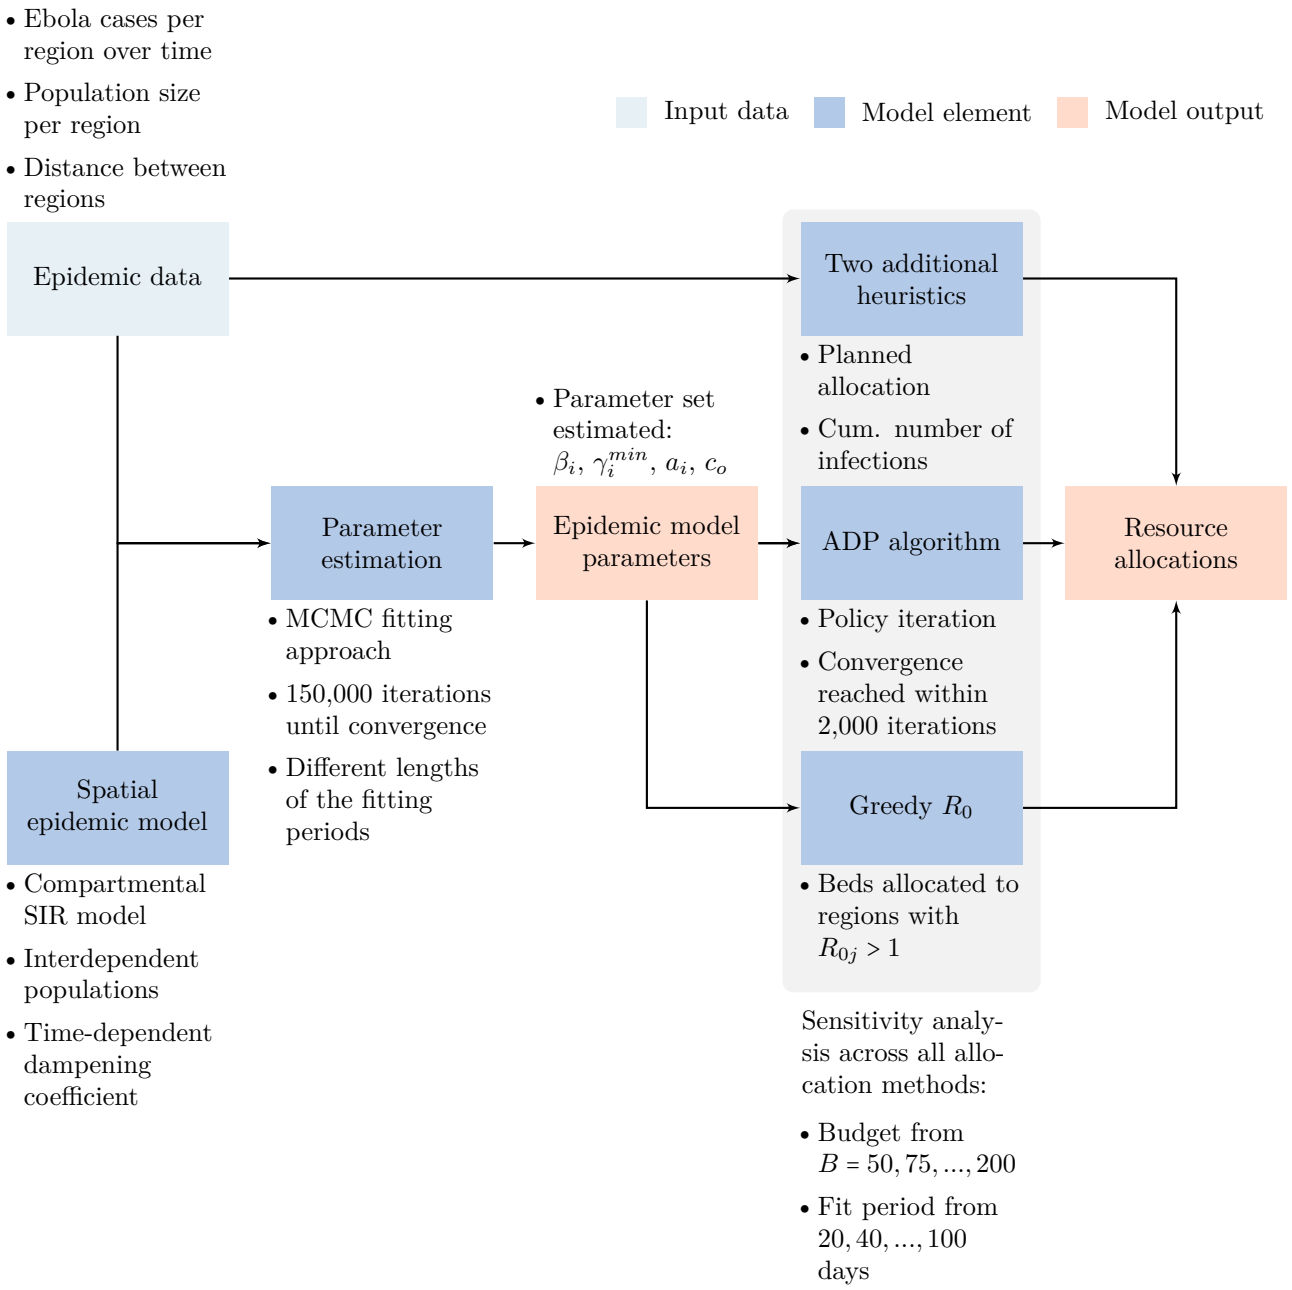

Figure C.1: Relationships between input data, model elements, and model results.

begin
  Set basis function ϕ_f;
  Set $\boldsymbol{\theta}_t^0$ for all t;
  Set starting state x_0;
for $n = 1, 2, \ldots, N$ **do**
  Draw randomly $e \sim \text{Uniform}(0, 1)$;
  for $t = t_0, t_0 + 1, \ldots, t_f$ **do**
    if $\sum_{i=1}^{K} \sum_{t_0}^{t} a_{i,t}^n \geq B$ **then**
      $a_{i,t}^n = 0 \ \forall i = 1, \ldots, K$;
    else if $(n < 200) \vee (e < 1/0.1n)$ **then**
      Randomly select $a_{i,t}^n \in \mathcal{A}_t$;
    else
      Compute $a_{i,t}^n = \underset{a_{i,t} \in \mathcal{A}_t}{\text{argmin}} \big(\sum_{i=1}^{K} I_{i,t}^{\text{new}}(x_t, a_{i,t}) + \gamma \boldsymbol{\theta}_t^{n-1} \boldsymbol{\phi}(x_{t+1}^n \mid x_t^n, a_{i,t})^T \big)$;
    end
  end
  Initialize $\hat{V}_{t_f}^n = 0$;
  for $t = t_f, t_f - 1, \ldots, t_0$ **do**
    $\hat{V}_t^n = \sum_{i=1}^{K} I_{i,t}^{\text{new}}(x_t^n, a_{i,t}^n) + \gamma \hat{V}_{t+1}^n$;
    $\boldsymbol{\theta}_t^n = \boldsymbol{\theta}_t^{n-1} - \frac{1}{\gamma^n} G^{n-1} \boldsymbol{\phi}^n \Delta^n$;
  end
end
Return basis function parameters $\boldsymbol{\theta}_t^N$ for all t.

Figure C.2: Approximate dynamic programming algorithm.

problems can be solved with an approximation method introduced by Hochbaum (1995), which is briefly explained in the following paragraph.

For each subperiod $\tau \in (0, 1, \ldots, r)$, we write the cost for allocating $a_{i,t+\tau}$ beds in

$t + \tau$ as:

$$c(a_{i,t+\tau}) = \sum_{i=1}^{K} I^{\text{new}}_{i,t+\tau}(x_{t+\tau}, a_{i,t+\tau}) + \gamma_r \theta^{n-1}_t \phi(x^n_{t+\tau+(1/r)} \mid x^n_{t+\tau}, a_{i,t+\tau})^T, \tag{C.1}$$

where γ_r is the subperiod discount rate. The approximate benefit p_k of allocating 1 bed in region k is $p_{k,\tau} = c(a_{k,\tau} = 0) - c(a_{k,\tau} = 1)$, where $a_i = 0 \ \forall i \neq k$. The resulting 0/1 Knapsack subproblem in period τ is thus:

$$\text{maximize} \quad \sum_{k=1}^{K} p_k a_{k,\tau} \tag{C.2}$$

$$\text{subject to} \quad \sum_{k=1}^{K} a_{k,\tau} \leq K, \tag{C.3}$$

$$a_{k,\tau} \in \{0, 1\} \quad \forall k. \tag{C.4}$$

After transforming the ADP algorithm's search for an optimal decision into a nonlinear Knapsack problem, we can solve large-scale resource allocation problems—with population sizes of several million—in a reasonable amount of time even though the underlying epidemic model is characterized by dynamics that are nonlinear and nonconvex.

C.2.2 Recursive least square updating

In line 19 of the ADP Algorithm, the parameter vector $\boldsymbol{\theta}^n_t$ is updated using a recursive least-squares approach for nonstationary data (Powell, 2011) that is based on the previous iteration's estimated coefficients $\boldsymbol{\theta}^{n-1}_t$:

$$\boldsymbol{\theta}^n_t = \boldsymbol{\theta}^{n-1}_t - \frac{1}{\gamma^n} G^{n-1} \boldsymbol{\phi}^n \Delta^n. \tag{C.5}$$

The $2K \times 2K$ matrix G^n is updated recursively using the expression

$$G^n = \frac{1}{\lambda^n}\left(G^{n-1} - \frac{1}{\gamma^n}(G^{n-1}\boldsymbol{\phi}^n(\boldsymbol{\phi}^n)^T G^{n-1})\right), \tag{C.6}$$

and γ^n is a scalar computed using

$$\gamma^n = \lambda^n + (\boldsymbol{\phi}^n)^T G^{n-1} \boldsymbol{\phi}^n. \tag{C.7}$$

The difference Δ^n between the previous iteration's value approximation for state x_t and its corresponding "observed" (i.e., simulated) value at iteration n is:

$$\Delta^n = \bar{V}_t^{n-1}(x_t) - \hat{V}_t^n(x_t). \tag{C.8}$$

C.2.3 Exploration versus Exploitation

One limitation of ADP is that an allocation decision (in the first loop of the algorithm) might visit only the states with a high approximated value. Previously unvisited, low-value states would be ignored even though their value approximation might change once visited. To address this problem, we explore the state space also by making random allocation decisions. Although exploration increases the likelihood of a near-optimal solution, it also reduces the speed of convergence—a trade-off well known as the "exploration–exploitation dilemma".

Heuristic learning policies (e.g., ϵ-greedy and Boltzmann exploration) or directed learning methods (e.g., knowledge gradient exploration (Powell, 2011) can determine an appropriate degree of exploration or a set of states for which exploration will likely be worthwhile. Because the exploration of our problem's state space is relatively inexpensive, we choose a combination of pure and ϵ-greedy exploration, which yielded good results in our experiments. During the first 100 iterations, the algorithm explores only by making random allocation decisions. Thereafter, during the nth iteration, if $e < 1/0.1n$ for $e \sim \text{Uniform}(0, 1)$ then a random decision is chosen; otherwise, the algorithm exploits the current value function approximation (Fig. C.3). This procedure assures both a high degree of exploration initially and decreasing exploration as the algorithm progresses.

C.3 Epidemic Model Calibration

The epidemic model was calibrated with confirmed Ebola case data from 21 regions in Guinea, Liberia, and Sierra Leone, which collectively accounted for 60% of the total infections observed (Humanitarian Data Exchange, 2015b). Regions with no change in cases, or with fewer than 50 cases or five data points, were omitted. We also excluded 61 observations (amounting to less than 2.5% of the total dataset) that

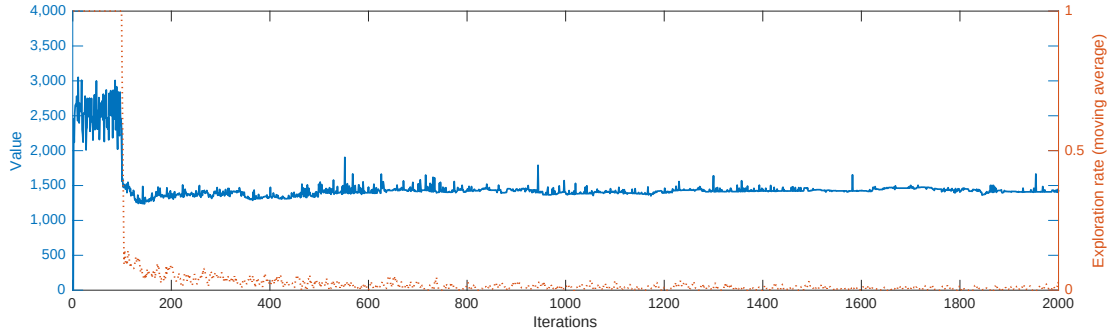

Figure C.3: Approximate dynamic program convergence (solid blue plot) and exploration rate (dotted red plot) with a budget of $B = 150$ beds, a fitting period of 60 days, and a maximum deployment rate of $r = 5$ beds per week per region.

could clearly be identified as outliers as well as two regions with inconsistent case numbers. These particular observations were likely the result of misreporting (e.g., in a single week one region's case count jumped from 50 to 1,000 before falling back to 60). We used 140 days of available data; of these, the first 100 days (beginning on 2 September 2014) were used for fitting purposes and the next 40 days were used to test the model's prediction against the outbreak's actual trajectory. As a robustness check, we evaluated the model's performance using case count data from only 20, 40, 60, 80, or 100 days. Model calibration was performed using a Markov chain Monte Carlo approach with a Metropolis–Hastings algorithm, an established method for fitting epidemic models to empirical data (Gibson, 1998; Currie, 2007). The model was simultaneously calibrated for all K regions in our dataset while assuming that the epidemic dynamics described in Equations (4.3)–(4.5). applied to each region. Because the data points were recorded on a daily basis, the model time horizon was divided into corresponding discrete periods: $t \in (t_0, \ldots, t_f)$.

The vector $\mathbf{w}_n$ contains the parameter c_0 and the parameters β_i, $\psi_i^{\max}$, and α_i, for $i = 1, 2, \ldots, K$, at iteration n of the MCMC algorithm. Since the value of each parameter is unknown, we assume Uniform *prior* distributions over the following intervals: $\beta_i \in (0, 10^{-5})$, $\psi_i^{\max} \in (0, 1)$, $\alpha_i \in (0, 1)$, and $c_0 \in (0.8, 1)$. The sum of squared errors between the observed Ebola case count $(\tilde{I}_{i,t})$ and the model projection $(I_{i,t}(\mathbf{w}_n))$ in region i at time t is:

$$\text{SSE} = \sum_{i=1}^{K} \sum_{t=t_0}^{t_f} \left(\tilde{I}_{i,t} - I_{i,t}(\mathbf{w}_n) \right)^2. \tag{C.9}$$

The Metropolis–Hastings algorithm recursively calculates a new parameter candidate

set $\mathbf{w}_n$ at iteration n via maximum likelihood with the following likelihood ratio:

$$\mathrm{LR} = \frac{P(\tilde{I}|\mathbf{w}_n)}{P(\tilde{I}|\mathbf{w}_{n-1})} = \exp\left(-\sum_{i=1}^{K}\sum_{t=t_0}^{t_f}\left(\tilde{I}_{i,t} - I_{i,t}(\mathbf{w}_n)\right)^2 + \sum_{i=1}^{K}\sum_{t=t_0}^{t_f}\left(\tilde{I}_{i,t} - I_{i,t}(\mathbf{w}_{n-1})\right)\right). \quad \text{(C.10)}$$

At iteration n, the candidate parameter set $\mathbf{w}_n$ replaces the previous set $\mathbf{w}_{n-1}$ if $\mathrm{LR} > \nu$, with $\nu \sim \mathrm{Uniform}(0,1)$. The algorithm continues until convergence is obtained for all parameter values (at most 150,000 iterations in our experiments; see Fig. C.25 - C.45 for convergence plots).

C.3.1 Epidemic model fit results

The accuracy of the developed epidemic model when comparing it to the observed case data is tested with varying amount of data available for calibration. Ranging from 20 to 100 days, the calibration periods represent the low data availability in the initial phase of an epidemic as well as more progressed situations when sufficient data is available. Projected and observed trajectories of the 21 regions can be found in Figure C.4 - C.24. Results show that the ability to accurately project the epidemic's trajectory increases with the amount of data available for calibrating the model. Results from Conakry, Guinea, and Margibi, Liberia, additionally indicate that not only the amount of available data, but also it's quality is a major factor for the projection accuracy. In the case of Conakry, a 40 days long fitting period results into a an acceptable MAPE of 9%, as the case data has no major jumps and is reported frequently. The observed data for Margibi is likely corrupted by underreporting between September and October 2014, resulting in a significantly higher MAPE of 88%. The variance of fit accuracies for the different fit periods are presented in Table C.1 to C.5.

Table C.1: Observed cases, projected cases, and MAPE for each region. The underlying epidemic model uses 20 days of data for parameter fitting and projects the epidemic trajectory for another 40. The MAPE is only computed for the timeframce after the period used for fitting.

Regions		Cases October 31st 2014		MAPE (%)
ID	Name	Projected	Observed	
1	Bo	82	175	32
2	Bombali	297	571	33
3	Bomi	34	73	19
4	Bong	65	78	12
5	Conakry	139	220	29
6	Coyah	10	52	64
7	Guéckédou	240	253	4
8	Kailahun	600	554	11
9	Kambia	19	39	42
10	Kenema	684	493	48
11	Kerouane	31	87	38
12	Kono	19	38	31
13	Lofa	276	304	5
14	Macenta	347	550	28
15	Margibi	38	312	34
16	Montserrado	412	1496	31
17	Moyamba	84	96	18
18	Nimba	66	106	24
19	Nzerekore	20	97	62
20	Port Loko	283	557	33
21	Tonkolili	80	210	42
Total / average		3823	6361	31

Table C.2: Observed cases, projected cases, and MAPE for each region. The underlying epidemic model uses 40 days of data for parameter fitting and projects the epidemic trajectory for another 40. The MAPE is only computed for the timeframce after the period used for fitting.

Regions		Cases November 20th 2014		MAPE (%)
ID	Name	Projected	Observed	
1	Bo	125	229	34
2	Bombali	660	745	13
3	Bomi	35	135	50
4	Bong	89	103	10
5	Conakry	227	249	9
6	Coyah	42	69	35
7	Guéckédou	253	260	2
8	Kailahun	540	559	3
9	Kambia	35	59	27
10	Kenema	444	494	9
11	Kerouane	37	139	58
12	Kono	40	72	19
13	Lofa	292	325	6
14	Macenta	470	640	18
15	Margibi	43	344	75
16	Montserrado	380	1584	65
17	Moyamba	146	154	8
18	Nimba	90	110	16
19	Nzerekore	47	166	57
20	Port Loko	521	771	20
21	Tonkolili	188	320	30
Total / average		4706	7527	27

Table C.3: Observed cases, projected cases, and MAPE for each region. The underlying epidemic model uses 60 days of data for parameter fitting and projects the epidemic trajectory for another 40. The MAPE is only computed for the timeframce after the period used for fitting.

Regions		Cases December 10th 2014		MAPE (%)
ID	Name	Projected	Observed	
1	Bo	183	288	28
2	Bombali	1058	892	8
3	Bomi	50	137	59
4	Bong	86	135	24
5	Conakry	299	289	5
6	Coyah	79	117	12
7	Guéckédou	261	264	1
8	Kailahun	557	563	1
9	Kambia	55	88	30
10	Kenema	482	494	4
11	Kerouane	85	145	40
12	Kono	55	123	30
13	Lofa	309	326	6
14	Macenta	563	699	14
15	Margibi	271	383	29
16	Montserrado	1172	1676	31
17	Moyamba	140	171	19
18	Nimba	111	114	6
19	Nzerekore	108	193	37
20	Port Loko	747	981	16
21	Tonkolili	334	390	18
Total / average		7006	8468	20

Table C.4: Observed cases, projected cases, and MAPE for each region. The underlying epidemic model uses 80 days of data for parameter fitting and projects the epidemic trajectory for another 40. The MAPE is only computed for the timeframce after the period used for fitting.

Regions		Cases December 30th 2014		MAPE (%)
ID	Name	Projected	Observed	
1	Bo	230	306	24
2	Bombali	1173	965	9
3	Bomi	79	139	44
4	Bong	114	150	21
5	Conakry	343	353	6
6	Coyah	97	155	27
7	Guéckédou	267	269	0
8	Kailahun	567	565	0
9	Kambia	73	111	29
10	Kenema	504	496	1
11	Kerouane	138	160	16
12	Kono	68	197	47
13	Lofa	320	332	4
14	Macenta	628	715	11
15	Margibi	427	391	4
16	Montserrado	1712	1771	6
17	Moyamba	210	185	6
18	Nimba	127	116	4
19	Nzerekore	161	209	26
20	Port Loko	945	1192	16
21	Tonkolili	457	429	5
Total / average		8641	9206	15

Table C.5: Observed cases, projected cases, and MAPE for each region. The underlying epidemic model uses 100 days of data for parameter fitting and projects the epidemic trajectory for another 40. The MAPE is only computed for the timeframce after the period used for fitting.

Regions		Cases January 19th 2014		MAPE (%)
ID	Name	Projected	Observed	
1	Bo	275	314	16
2	Bombali	1216	979	12
3	Bomi	113	139	21
4	Bong	129	150	17
5	Conakry	378	383	3
6	Coyah	125	168	28
7	Guéckédou	270	269	1
8	Kailahun	571	565	0
9	Kambia	97	140	25
10	Kenema	512	498	2
11	Kerouane	169	160	4
12	Kono	103	235	47
13	Lofa	325	332	3
14	Macenta	672	715	7
15	Margibi	453	391	8
16	Montserrado	1831	1775	2
17	Moyamba	236	202	11
18	Nimba	128	116	6
19	Nzerekore	212	211	6
20	Port Loko	1188	1293	10
21	Tonkolili	553	445	12
Total / average		9555	9480	11

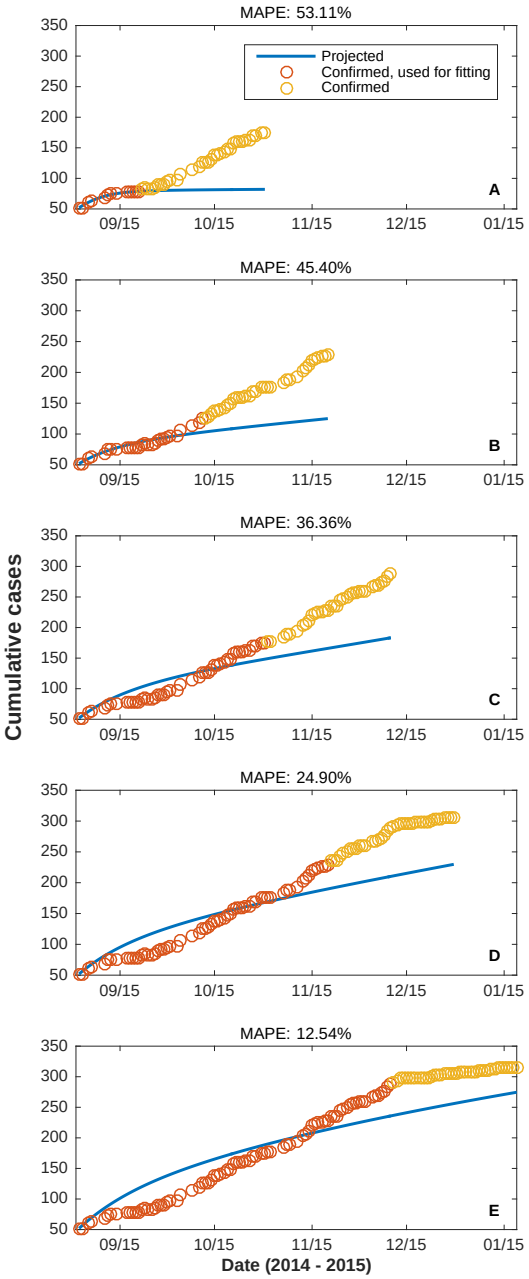

Figure C.4: Epidemic model fit for Bo, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

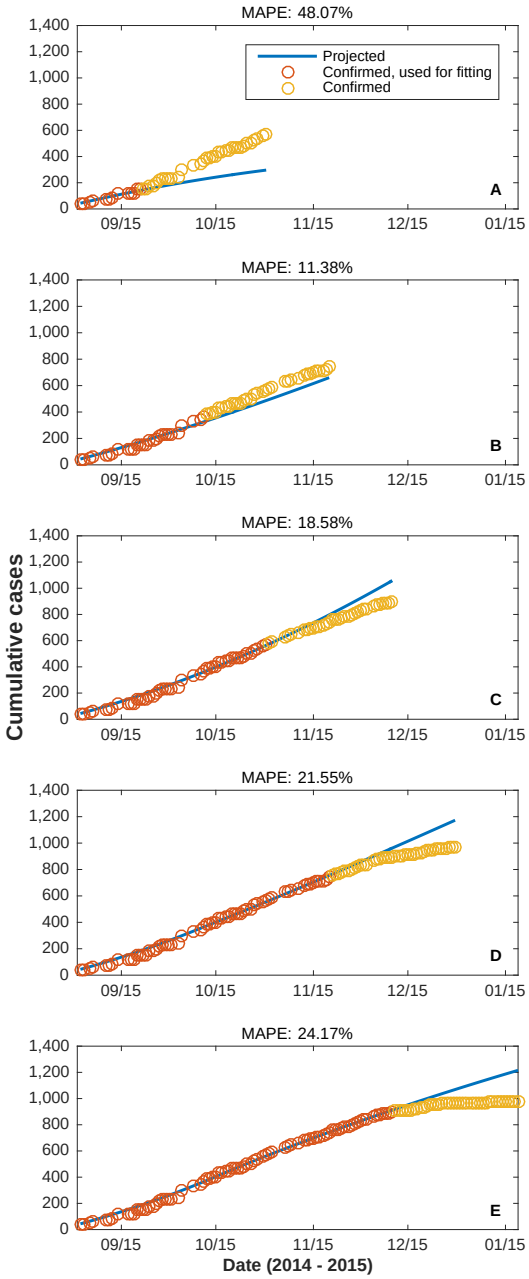

Figure C.5: Epidemic model fit for Bombali, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

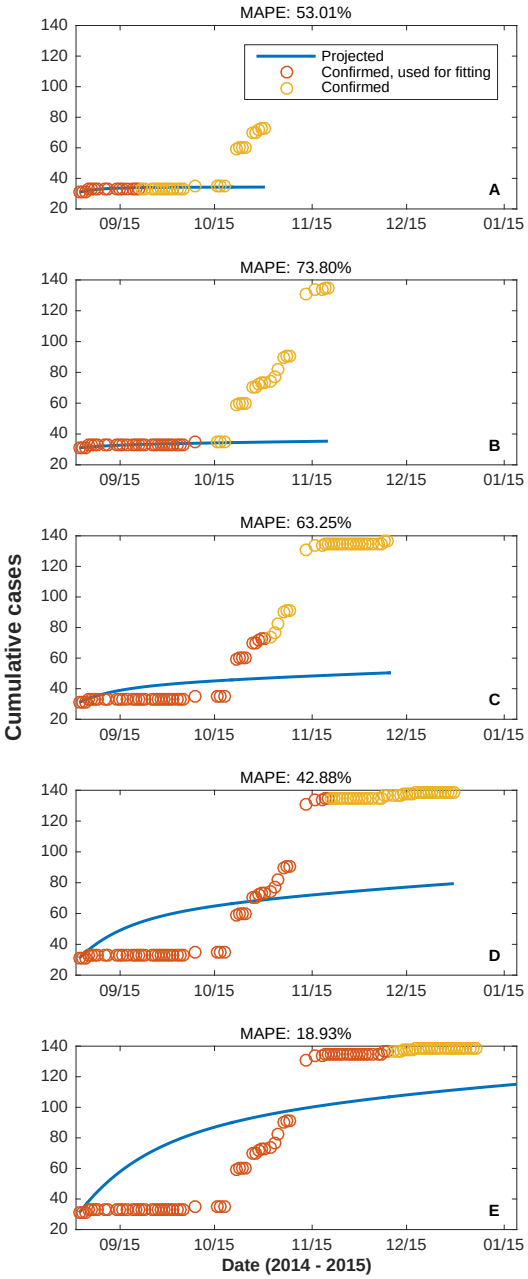

Figure C.6: Epidemic model fit for Bomi, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

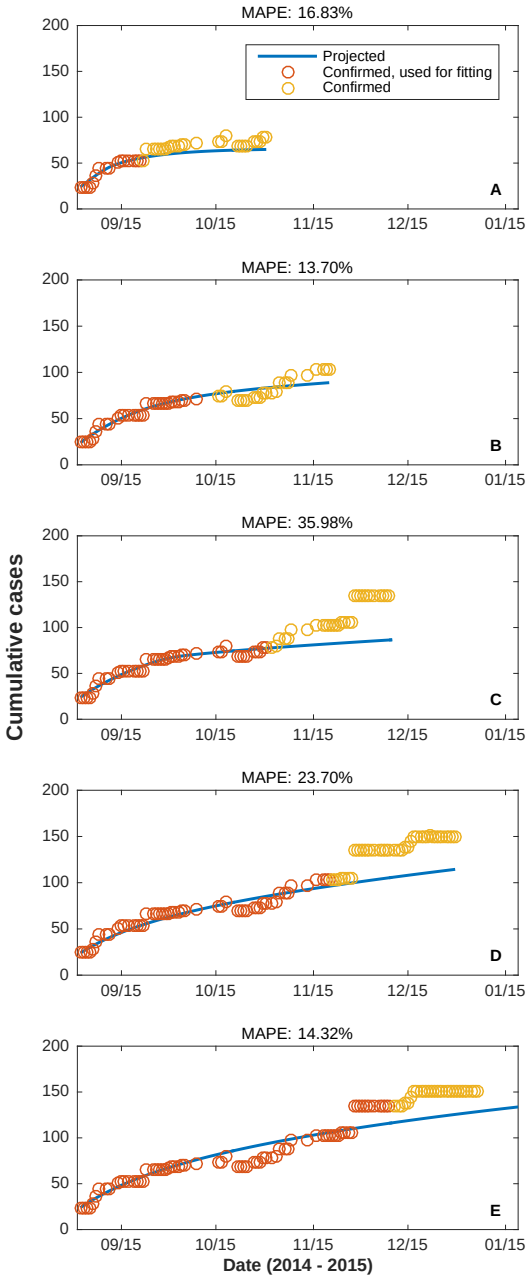

Figure C.7: Epidemic model fit for Bong, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

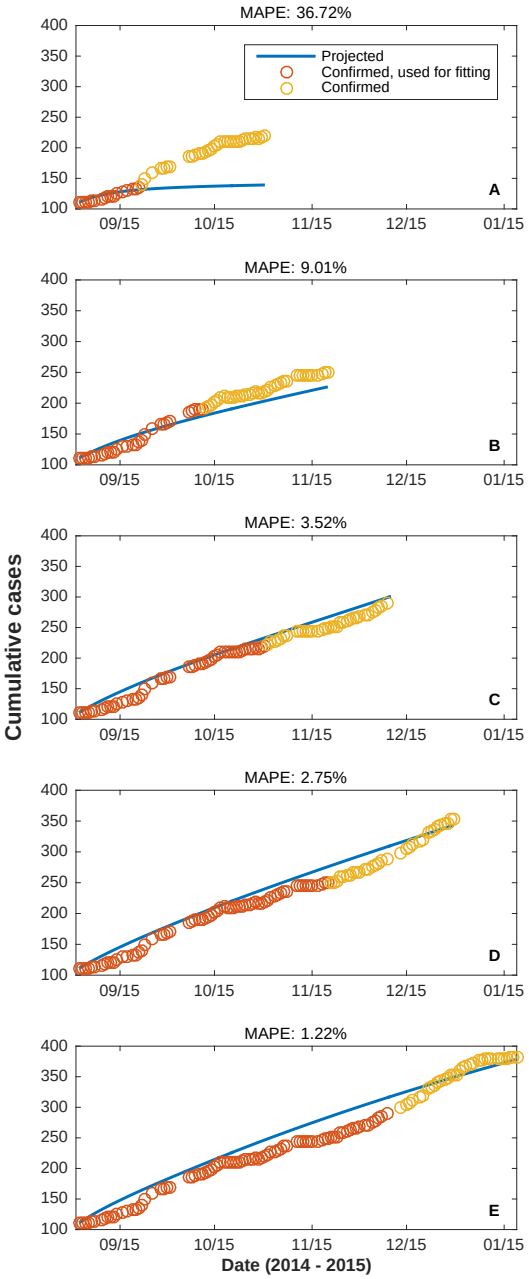

Figure C.8: Epidemic model fit for Conakry, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

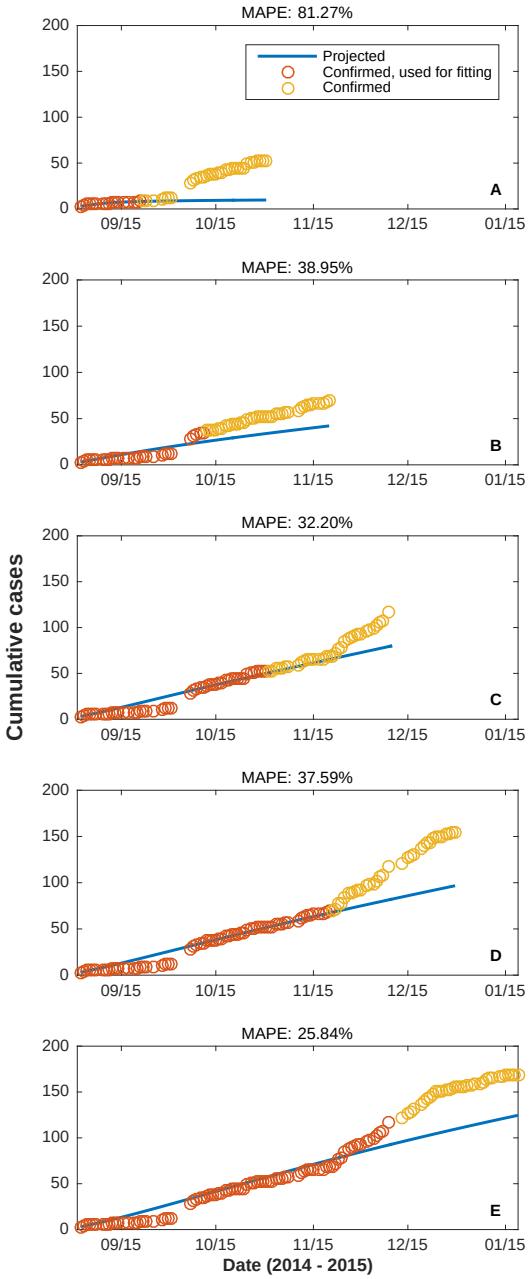

Figure C.9: Epidemic model fit for Coyah, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

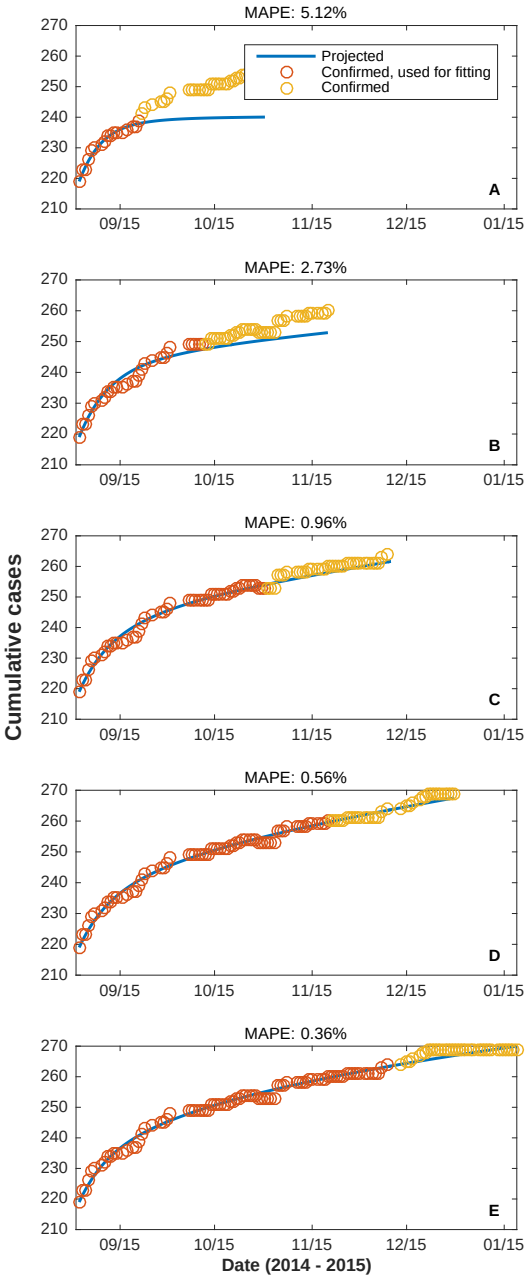

Figure C.10: Epidemic model fit for Guéckédou, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

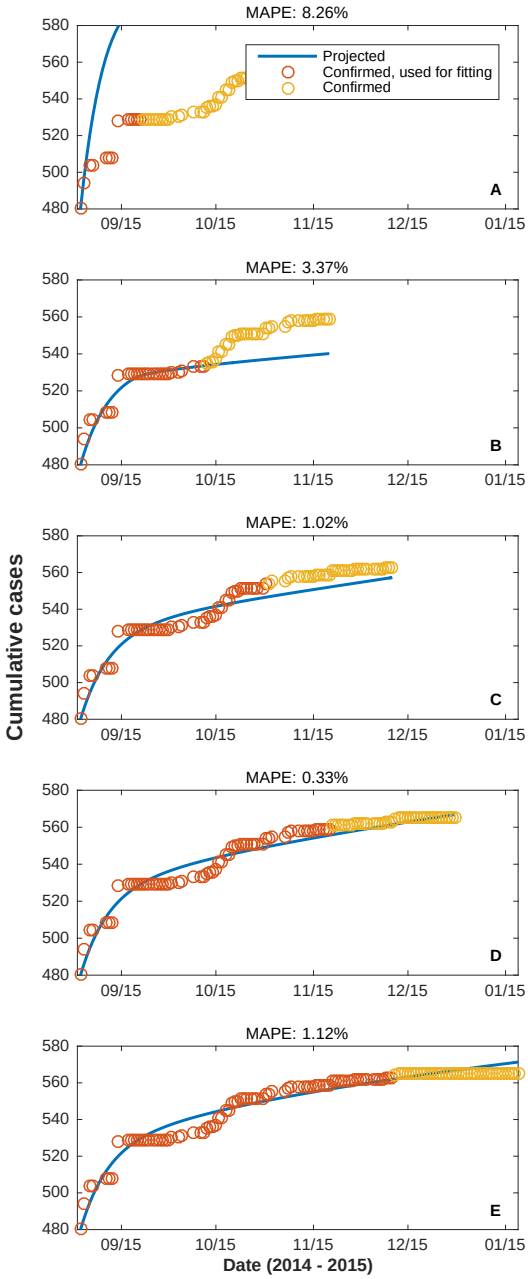

Figure C.11: Epidemic model fit for Kailahun, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

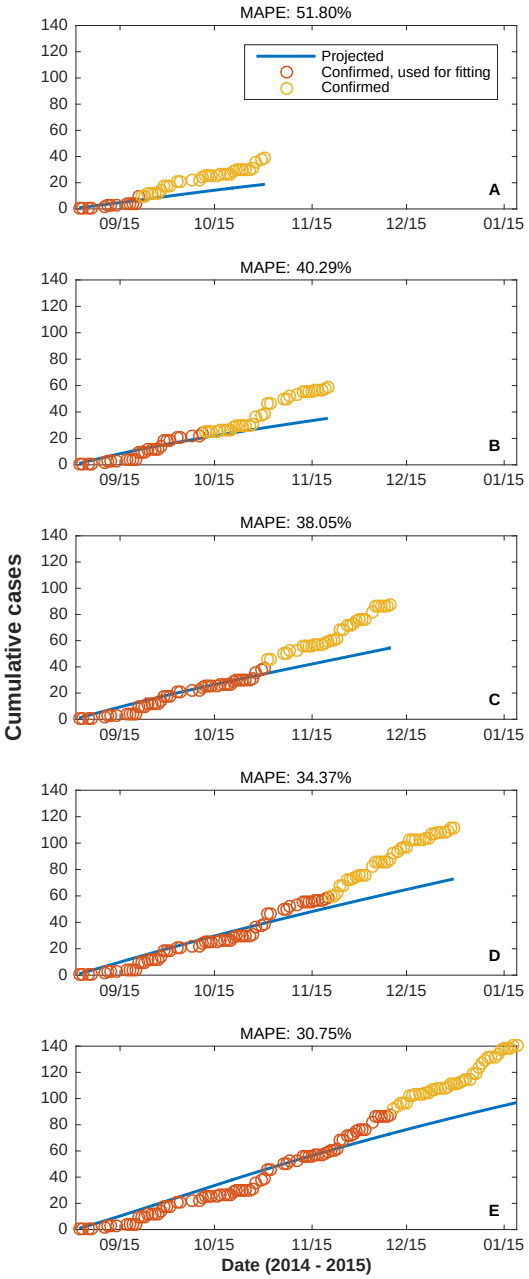

Figure C.12: Epidemic model fit for Kambia, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

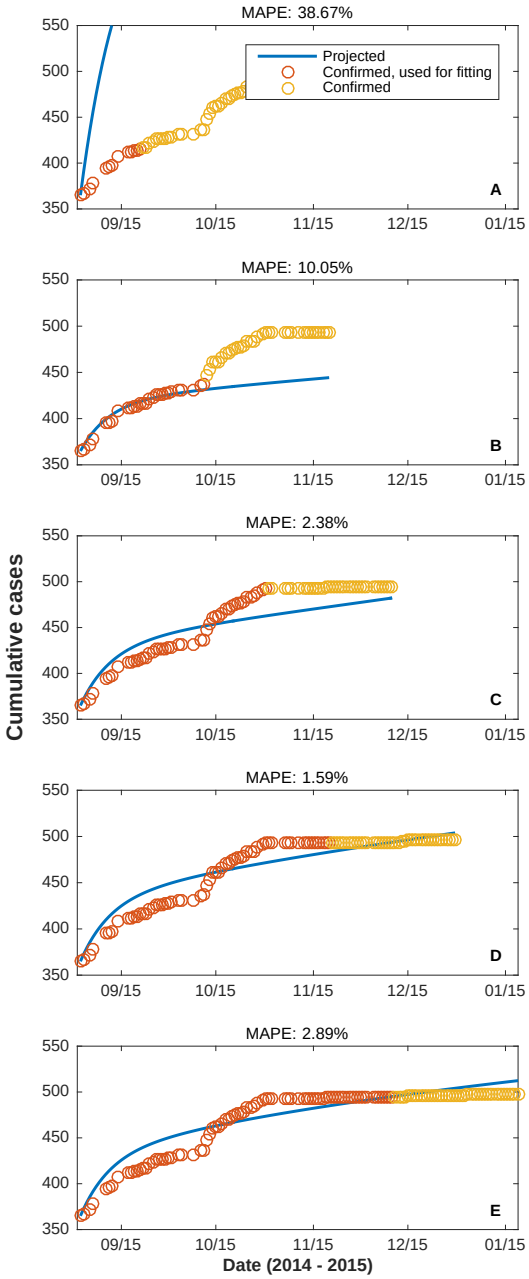

Figure C.13: Epidemic model fit for Kenema, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

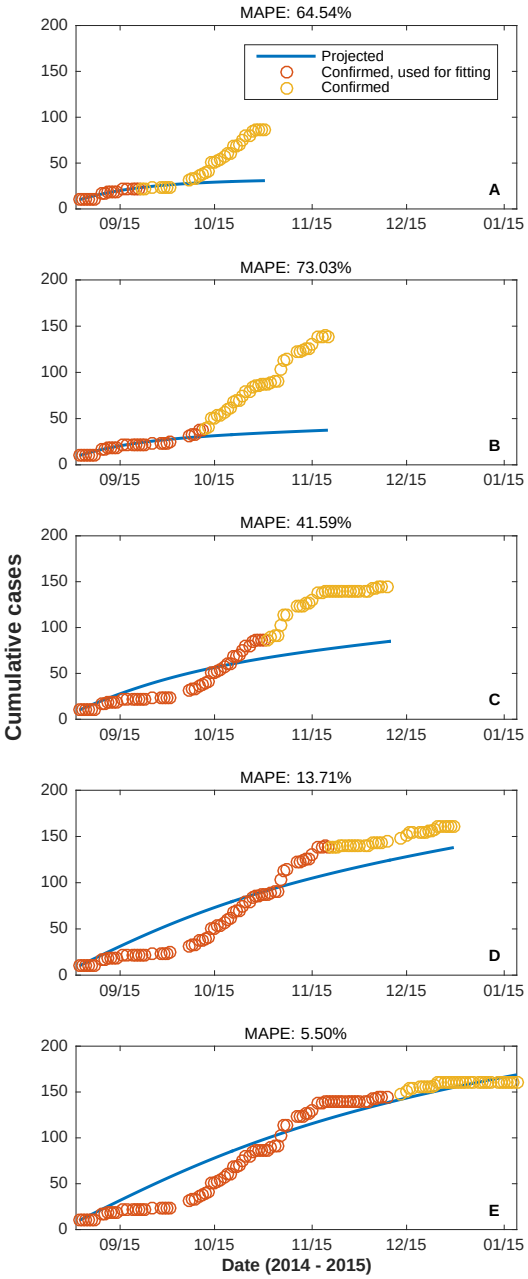

Figure C.14: Epidemic model fit for Kerouane, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

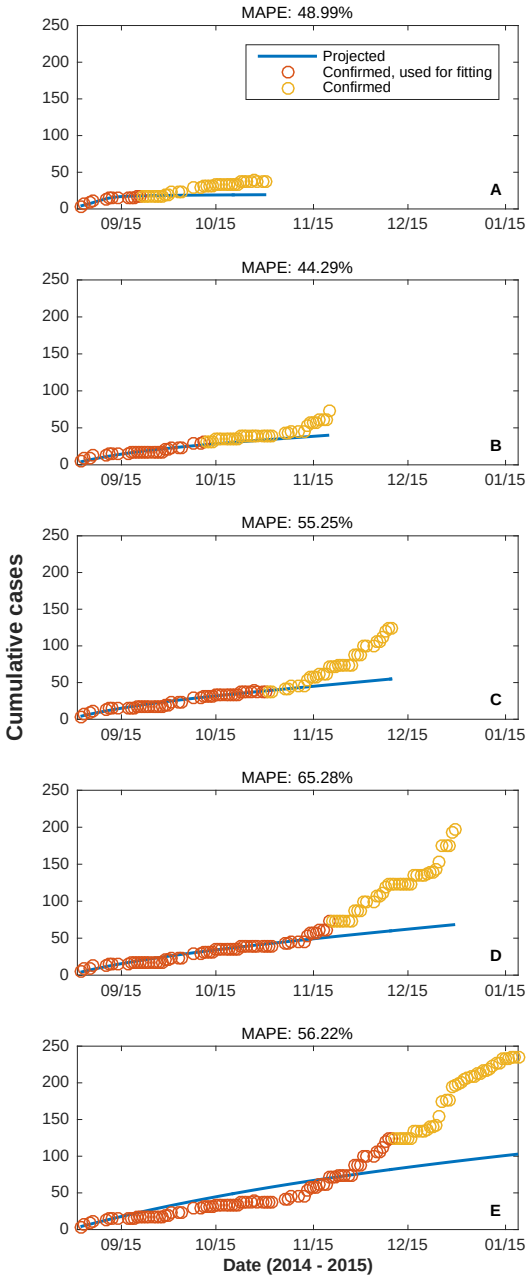

Figure C.15: Epidemic model fit for Kono, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

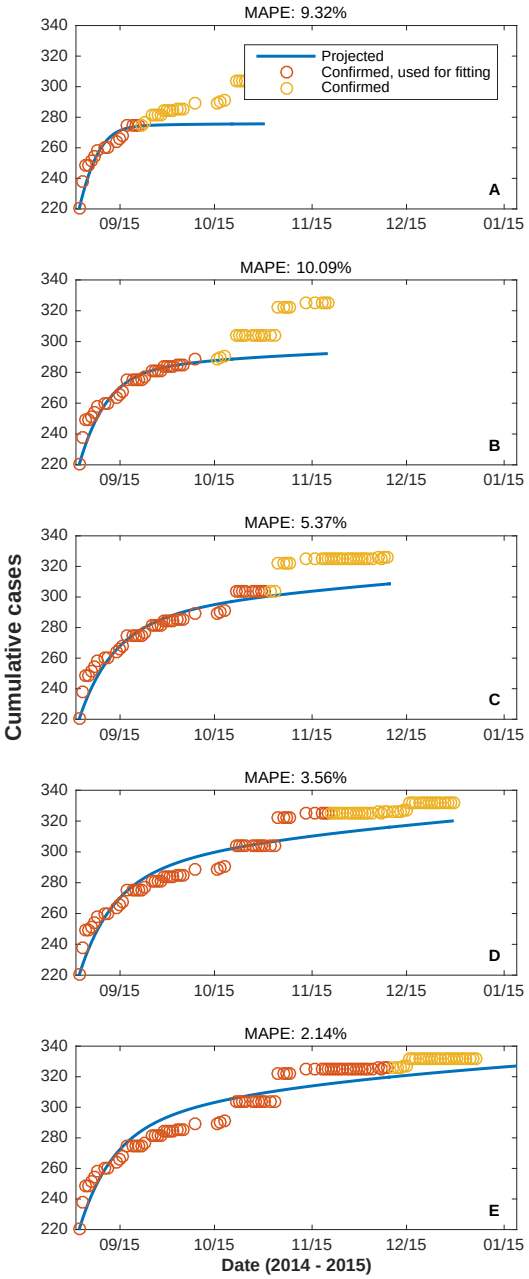

Figure C.16: Epidemic model fit for Lofa, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

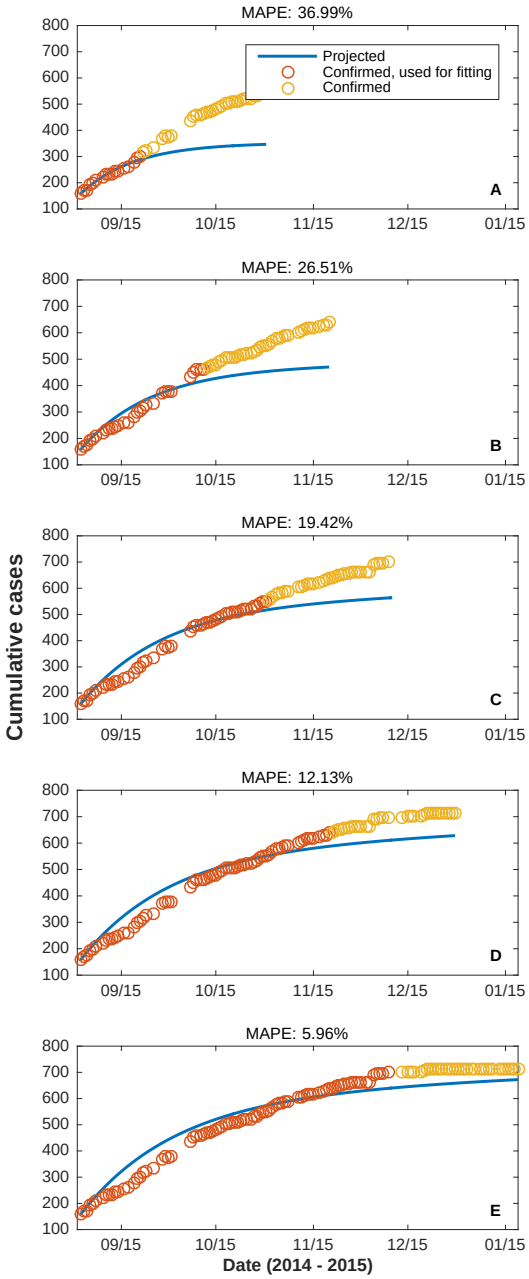

Figure C.17: Epidemic model fit for Macenta, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

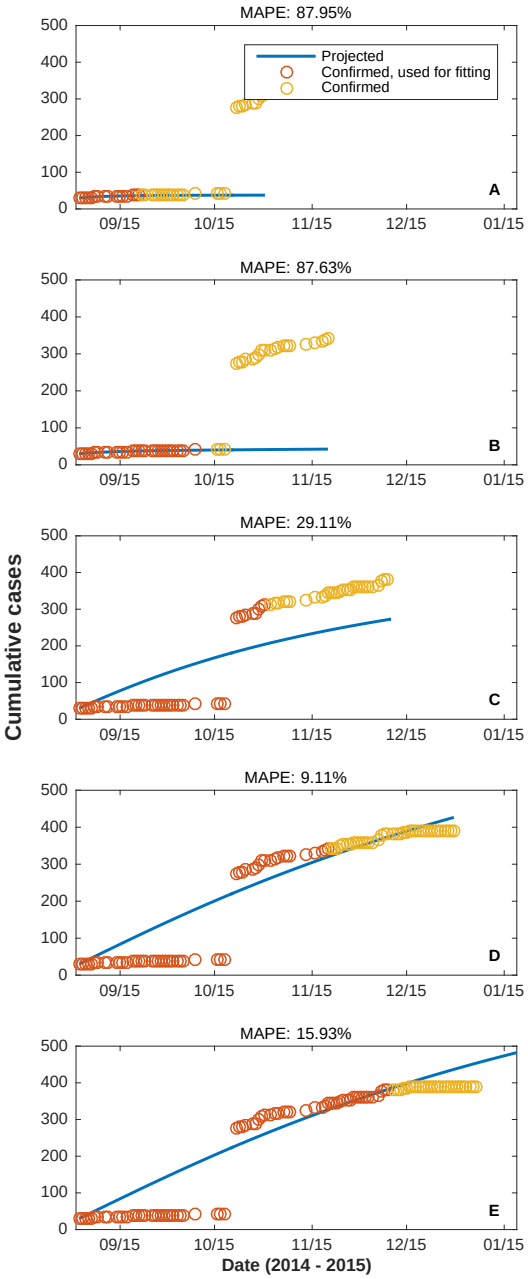

Figure C.18: Epidemic model fit for Margibi, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

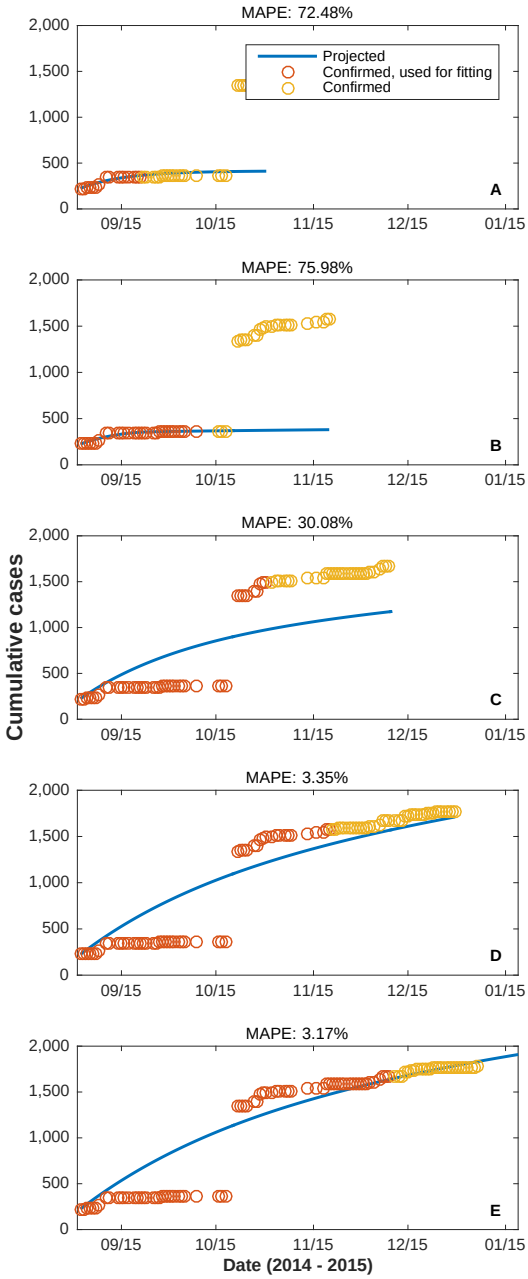

Figure C.19: Epidemic model fit for Montserrade, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

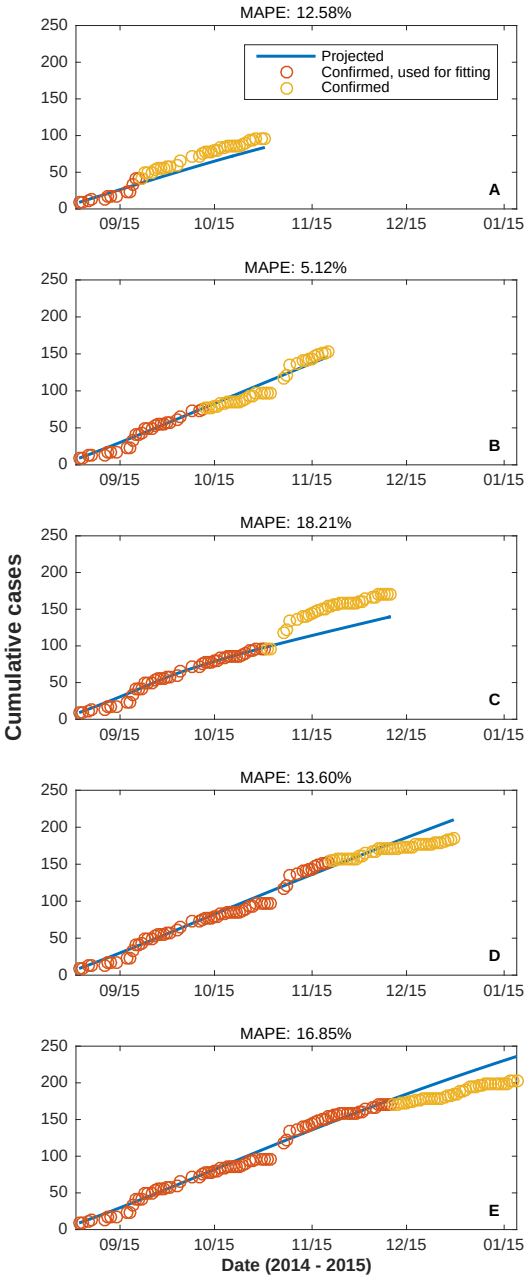

Figure C.20: Epidemic model fit for Moyamba, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

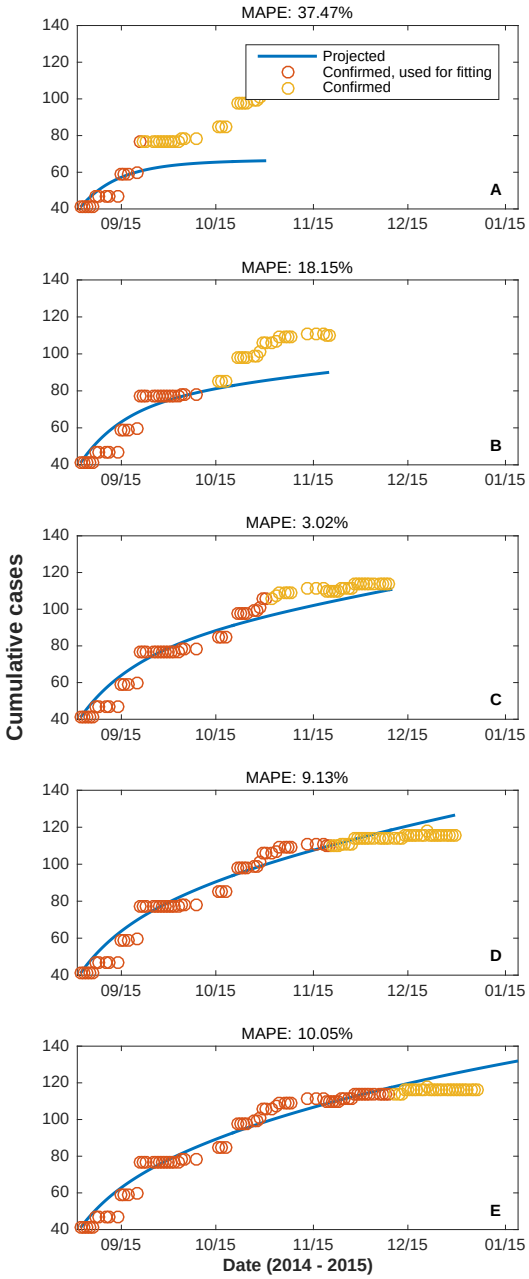

Figure C.21: Epidemic model fit for Nimba, Liberia, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

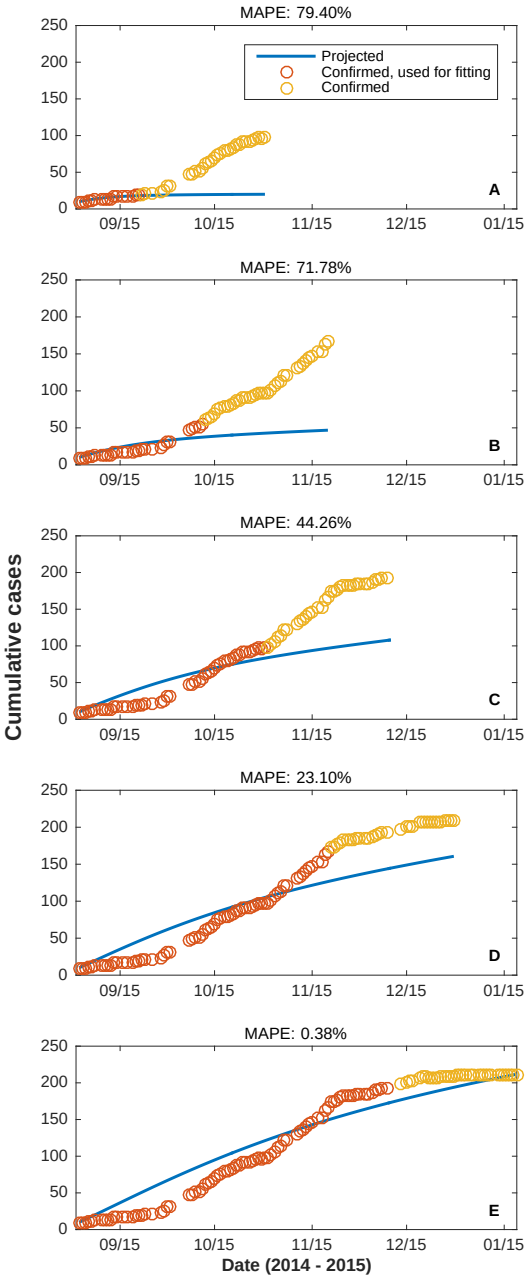

Figure C.22: Epidemic model fit for Nzereko, Guinea, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

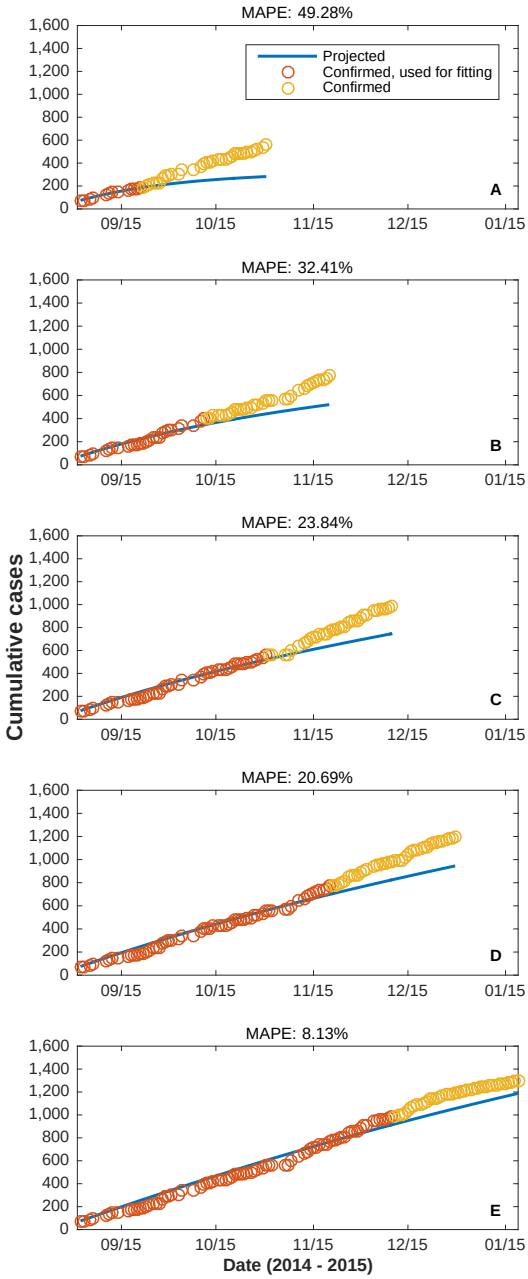

Figure C.23: Epidemic model fit for Porto Loko, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

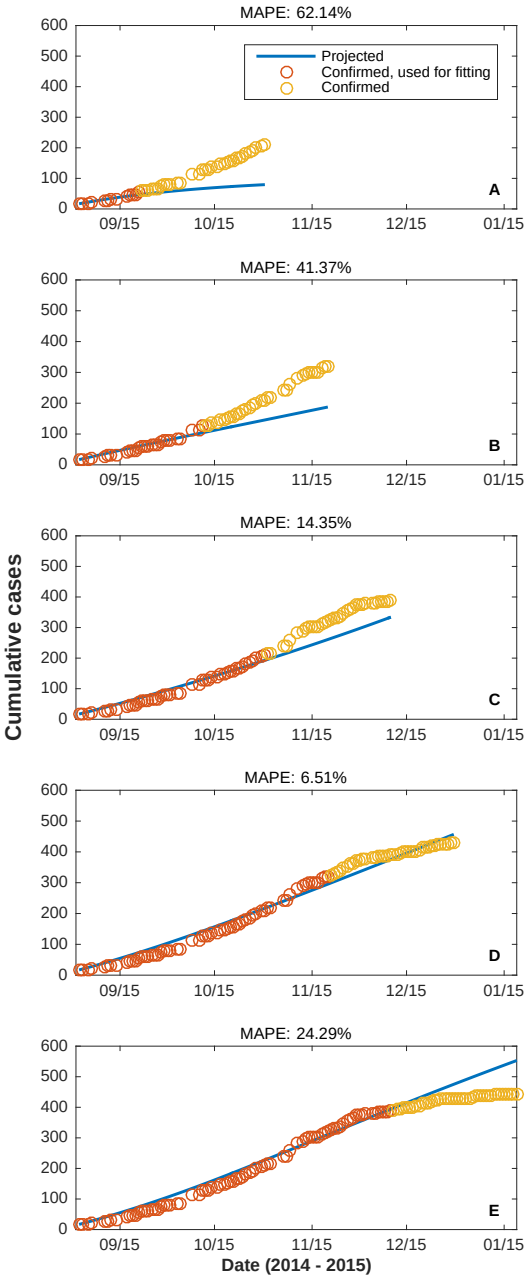

Figure C.24: Epidemic model fit for Tonkolili, Sierra Leone, with (A) 20, (B), 40, (C), 60, (D) 80 and, (E) 100 days of data used for fitting and a projection period of 40 days.

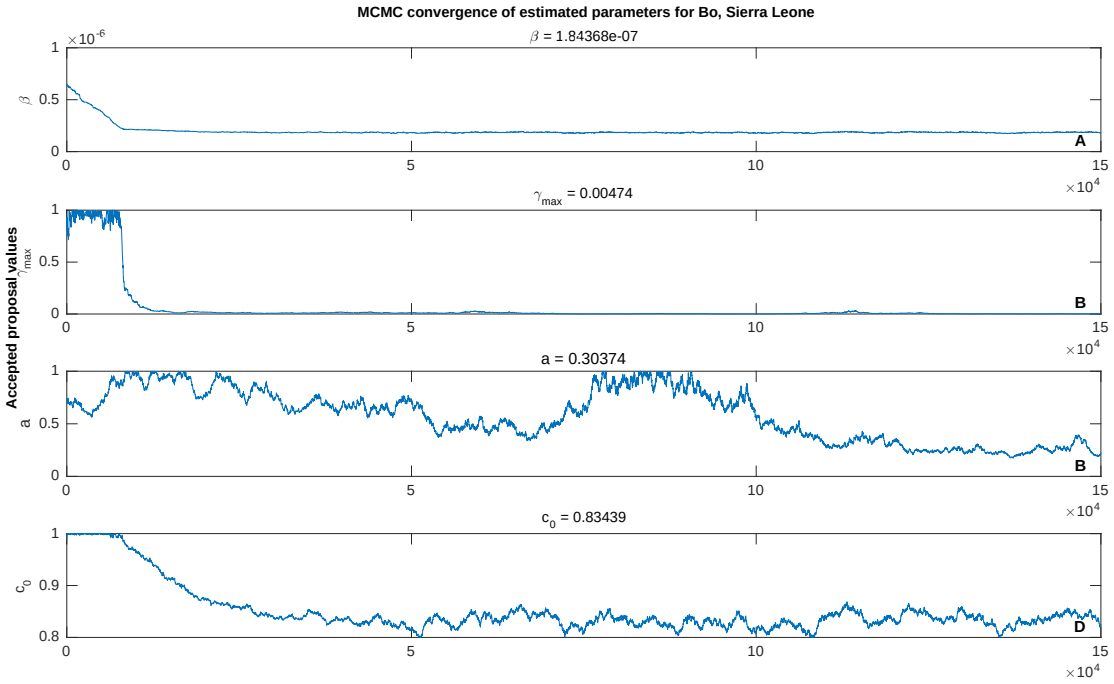

Figure C.25: MCMC convergence for Bo, Sierra Leone. The algorithm fits 4 parameters to 100 days of observed data: (A) The transmission coefficient β, (B) the maximum dampening factor γ^{max}, (C) the steepness factor of the damening curve a, and (D) the share of contacts inside a popoulation c_o. The panel description for the other regions is identical and therefore omitted in the remaining figures.

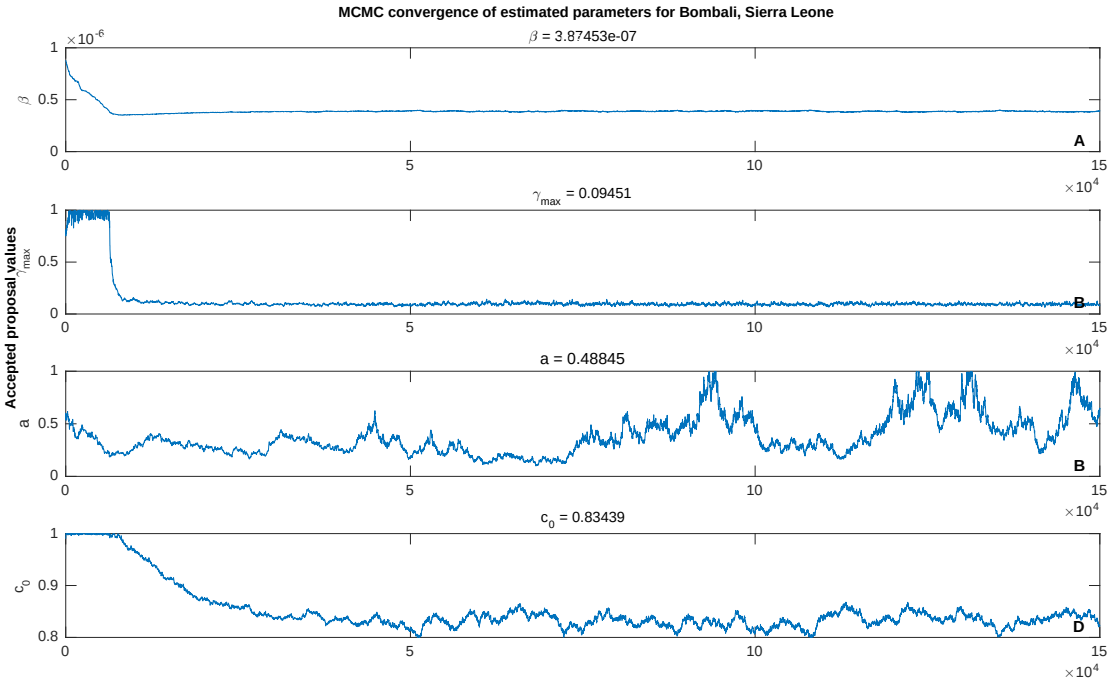

Figure C.26: MCMC convergence for Bombali, Sierra Leone.

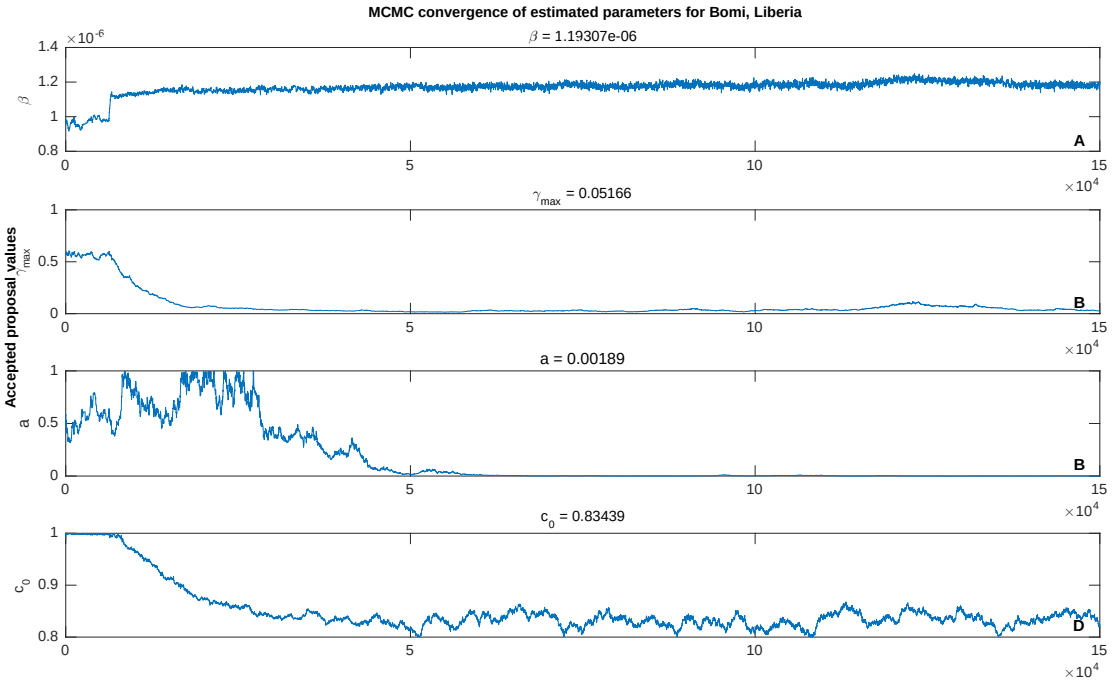

Figure C.27: MCMC convergence for Bomi, Liberia.

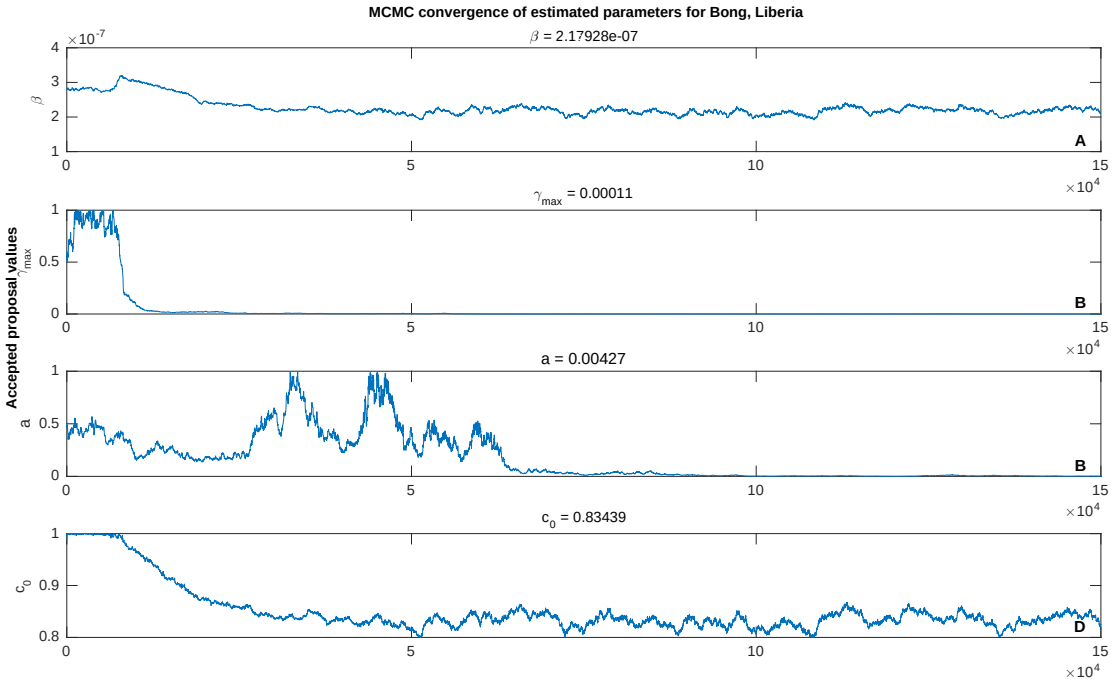

Figure C.28: MCMC convergence for Bong, Liberia.

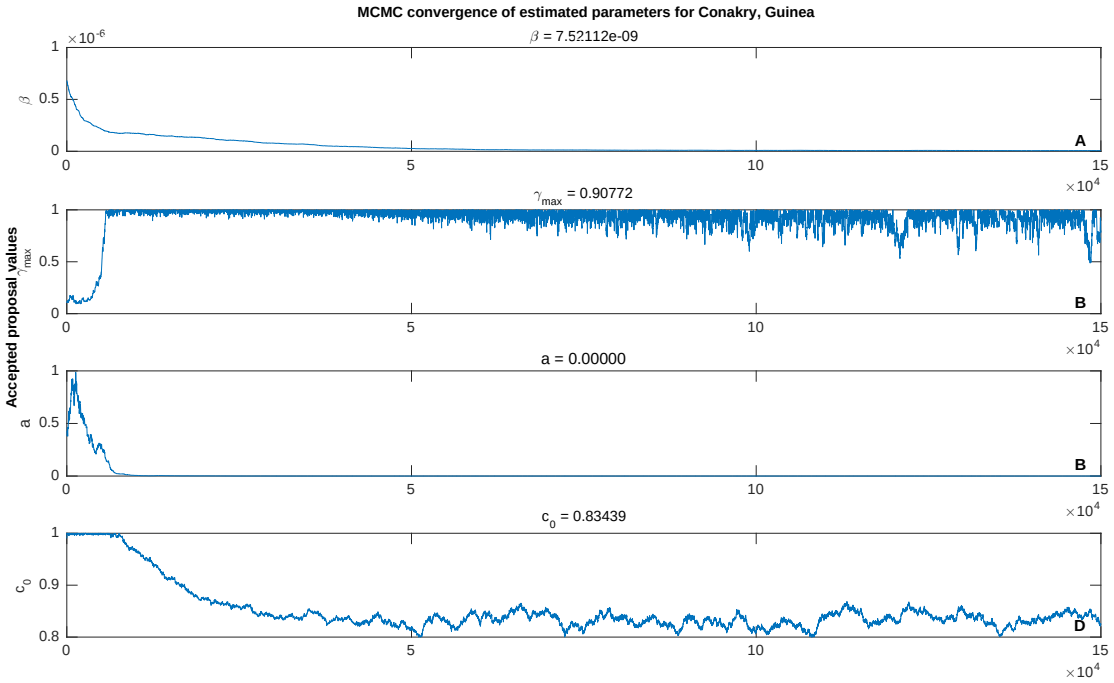

Figure C.29: MCMC convergence for Conakry, Guinea.

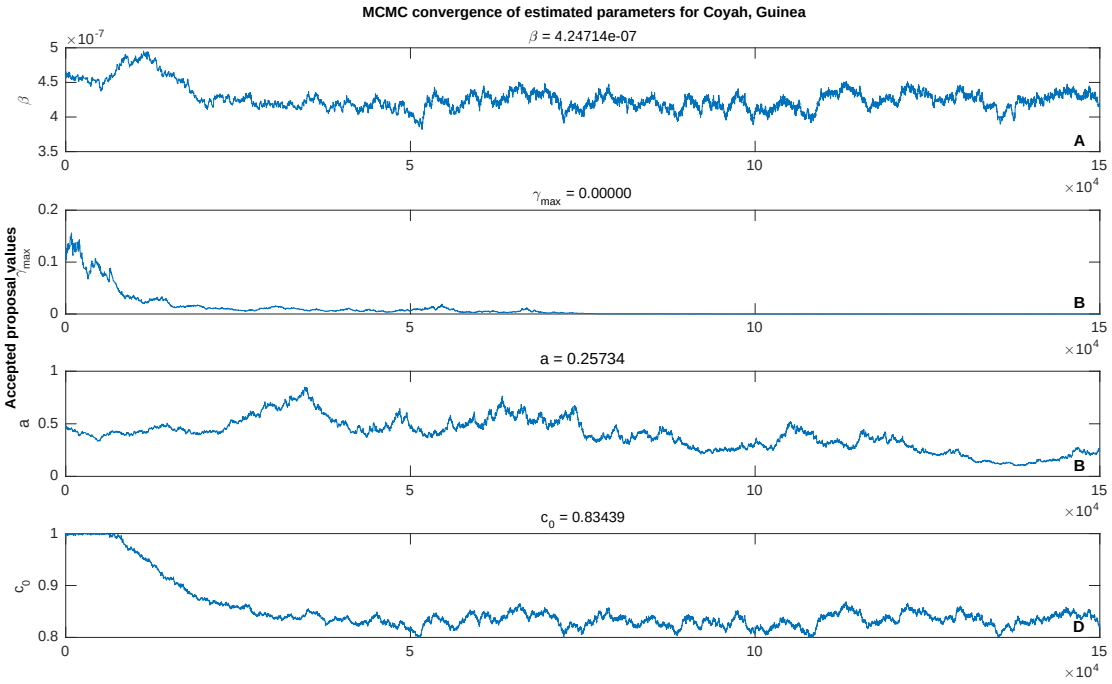

Figure C.30: MCMC convergence for Coyah, Guinea.

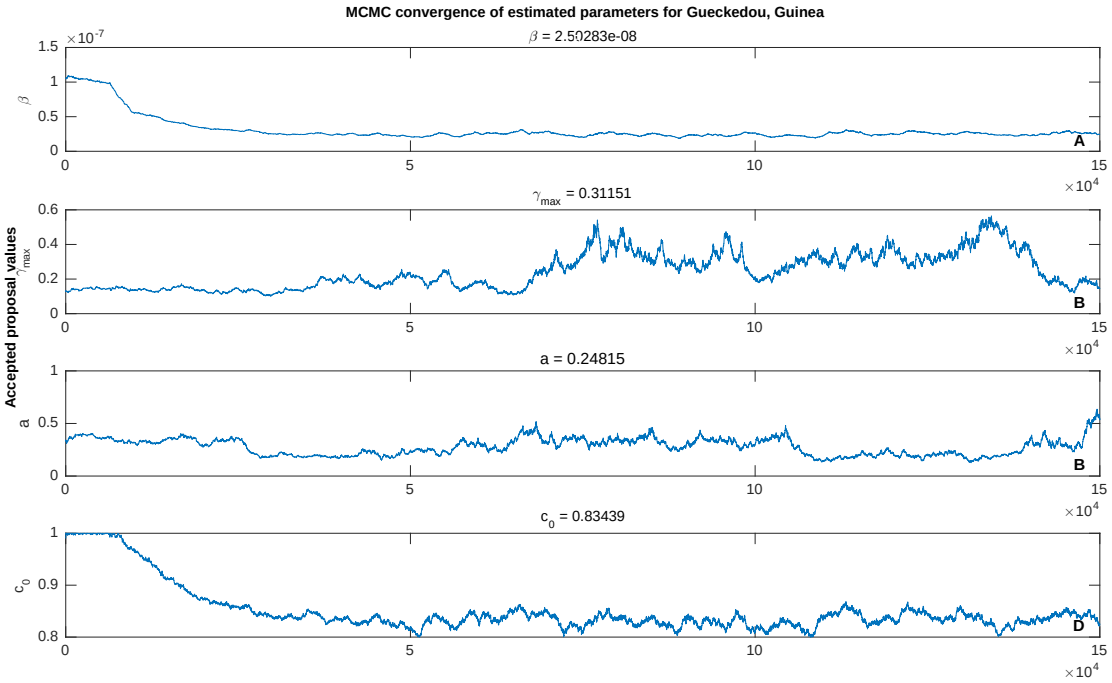

Figure C.31: MCMC convergence for Guéckédou, Guinea.

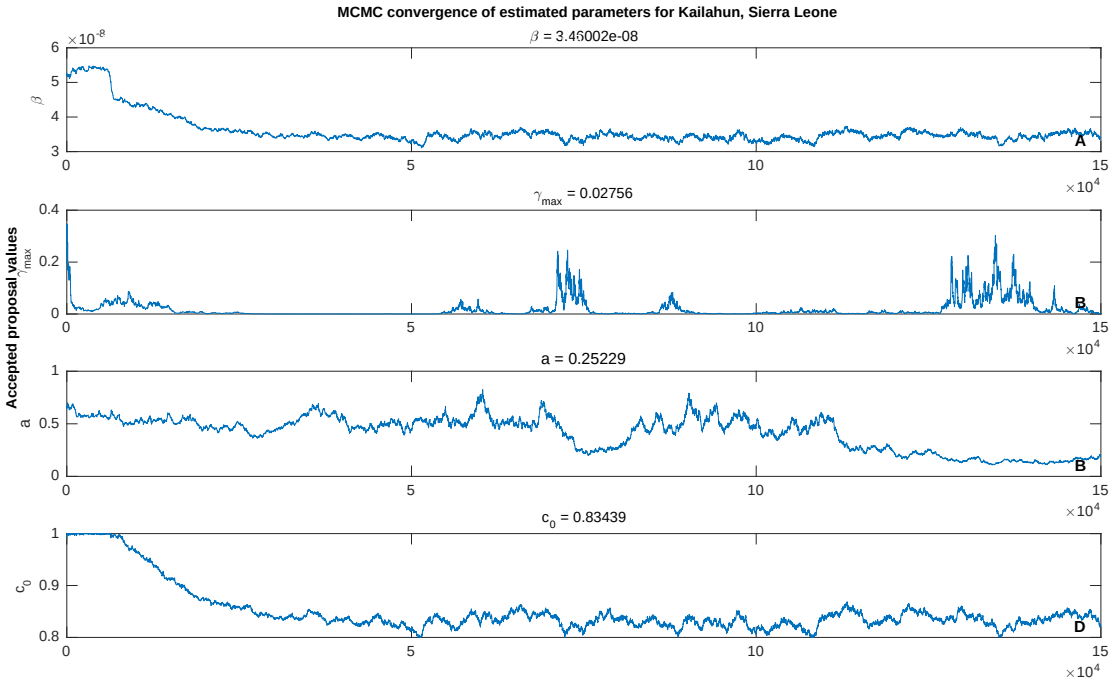

Figure C.32: MCMC convergence for Kailahun, Sierra Leone.

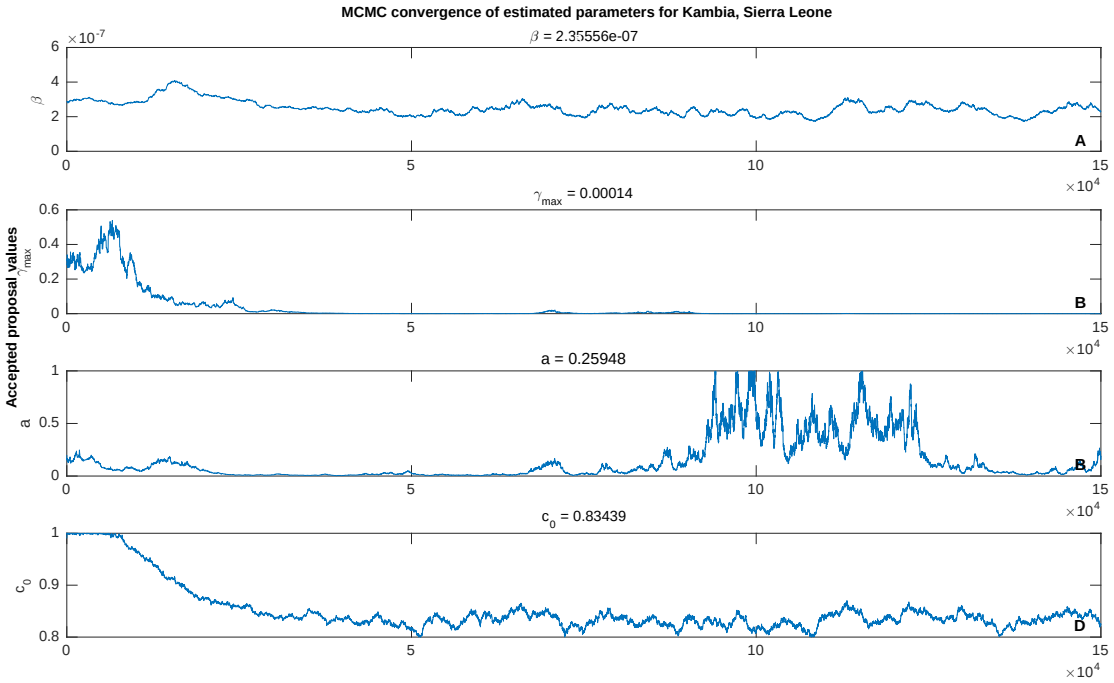

Figure C.33: MCMC convergence for Kambia, Sierra Leone.

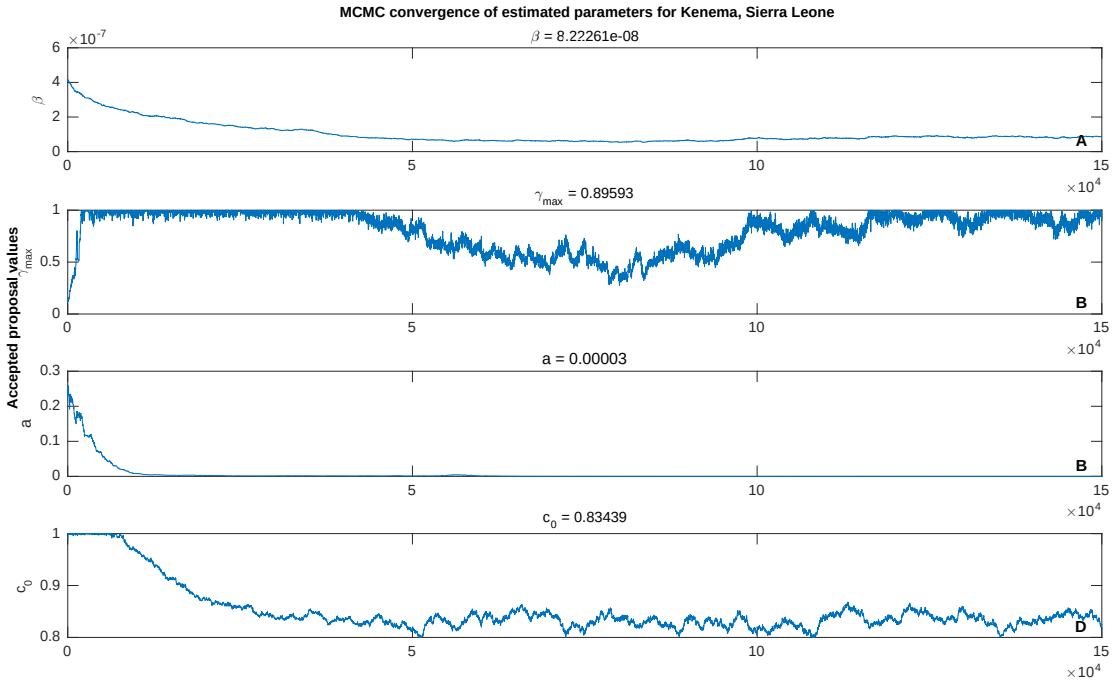

Figure C.34: MCMC convergence for Kenema, Sierra Leone.

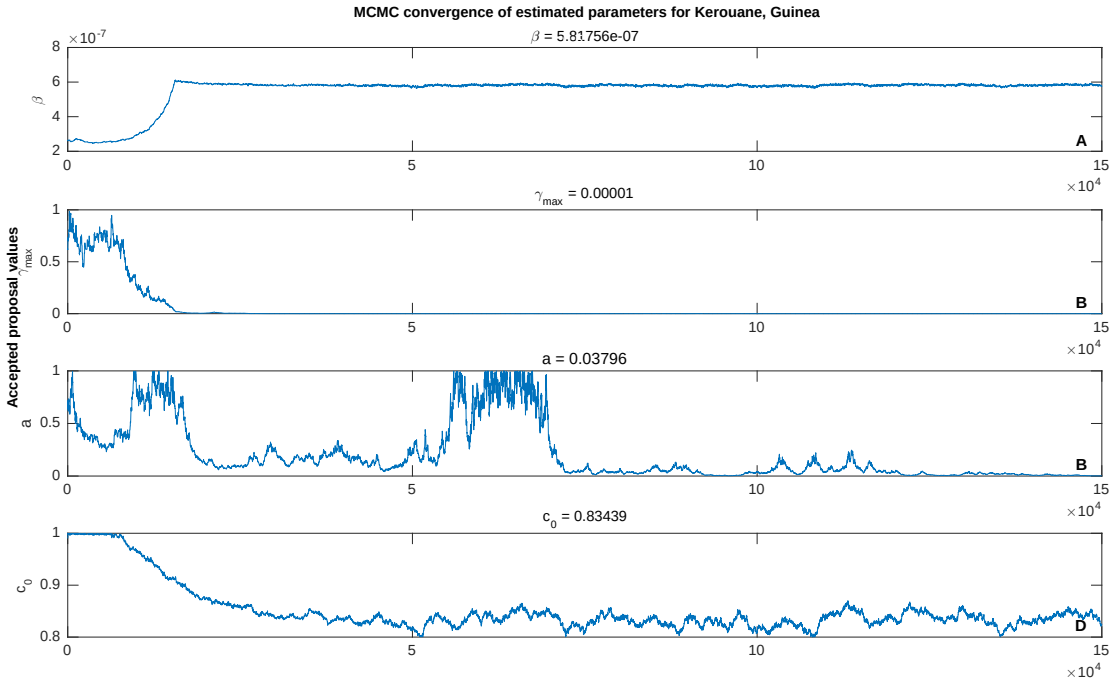

Figure C.35: MCMC convergence for Kerouane, Guinea.

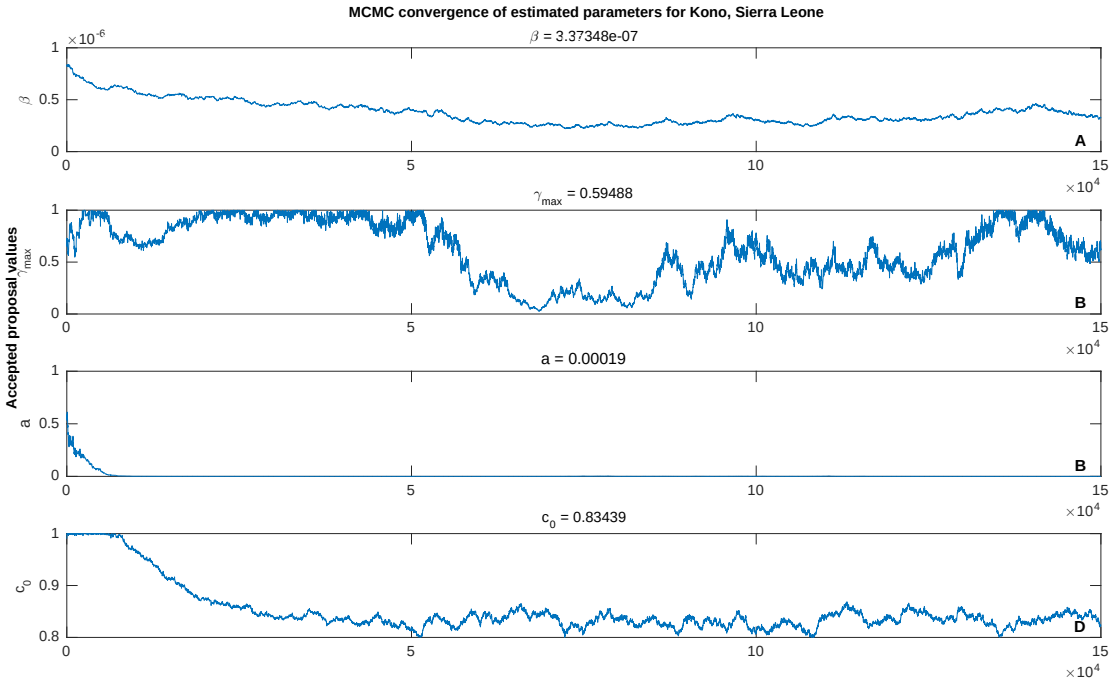

Figure C.36: MCMC convergence for Kono, Sierra Leone.

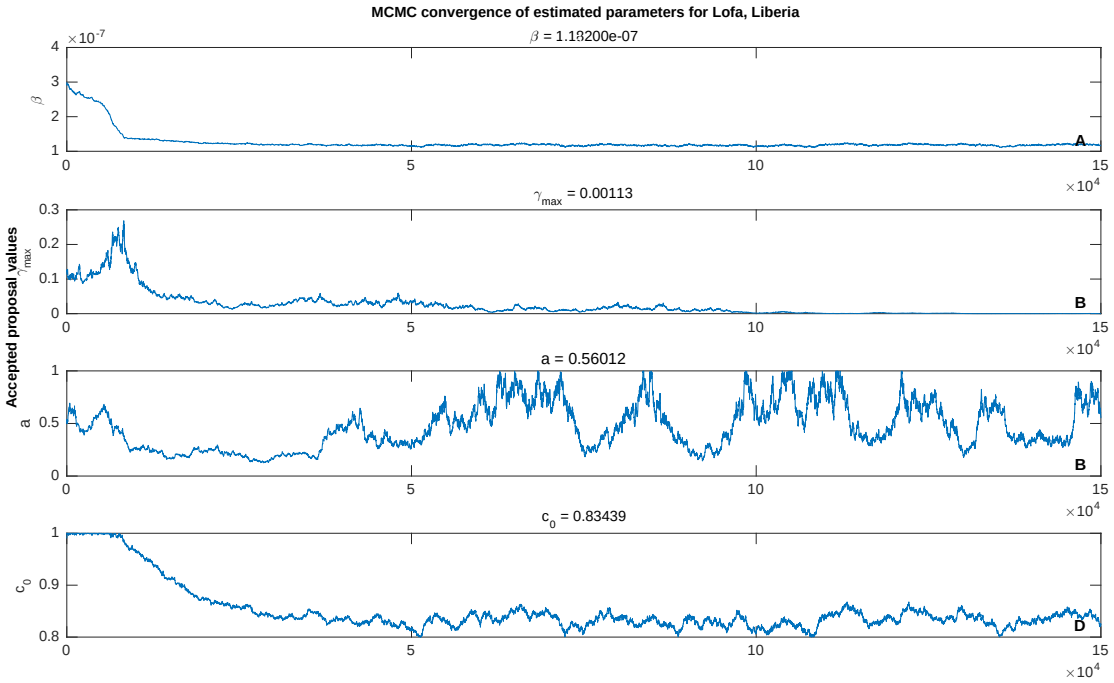

Figure C.37: MCMC convergence for Lofa, Liberia.

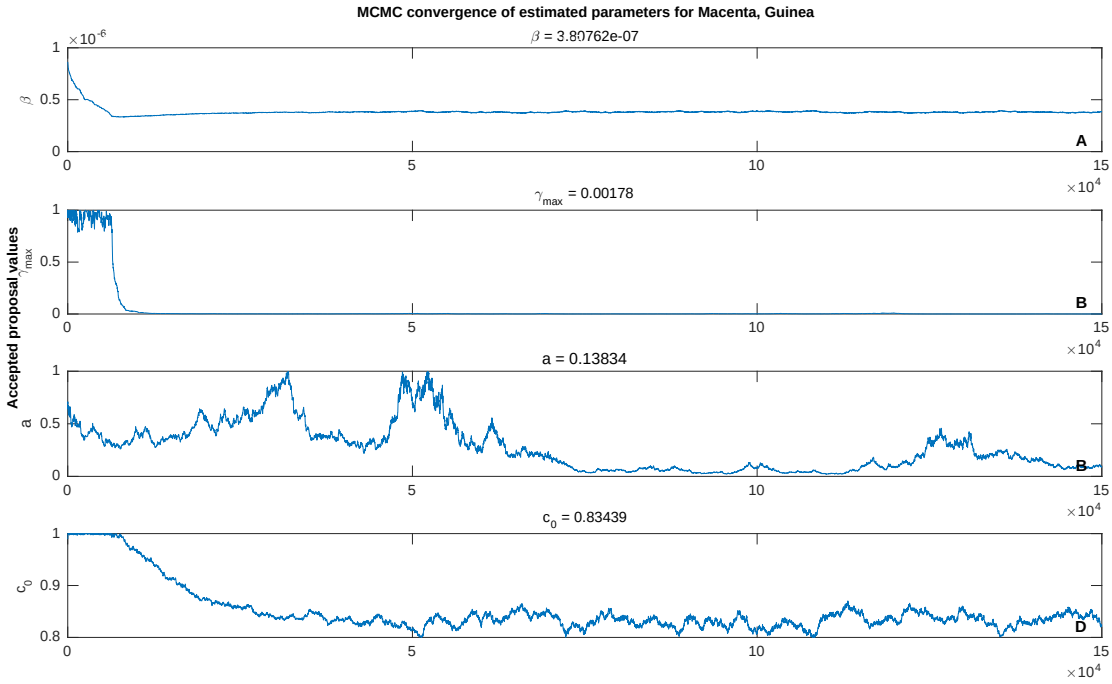

Figure C.38: MCMC convergence for Macenta, Guinea.

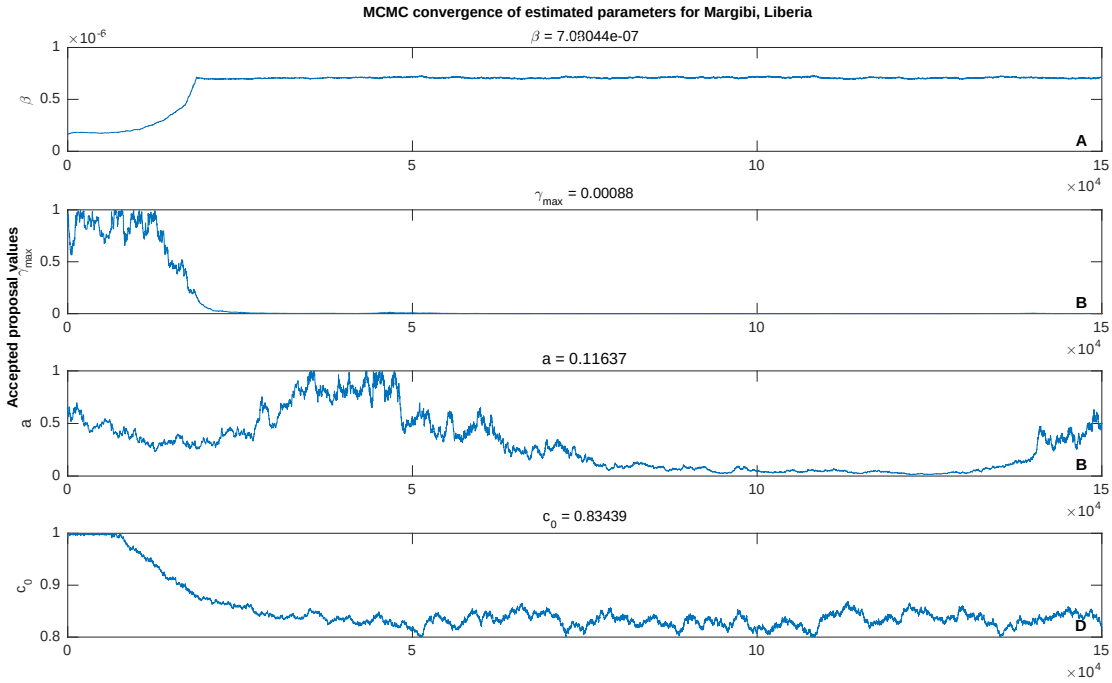

Figure C.39: MCMC convergence for Margibi, Liberia.

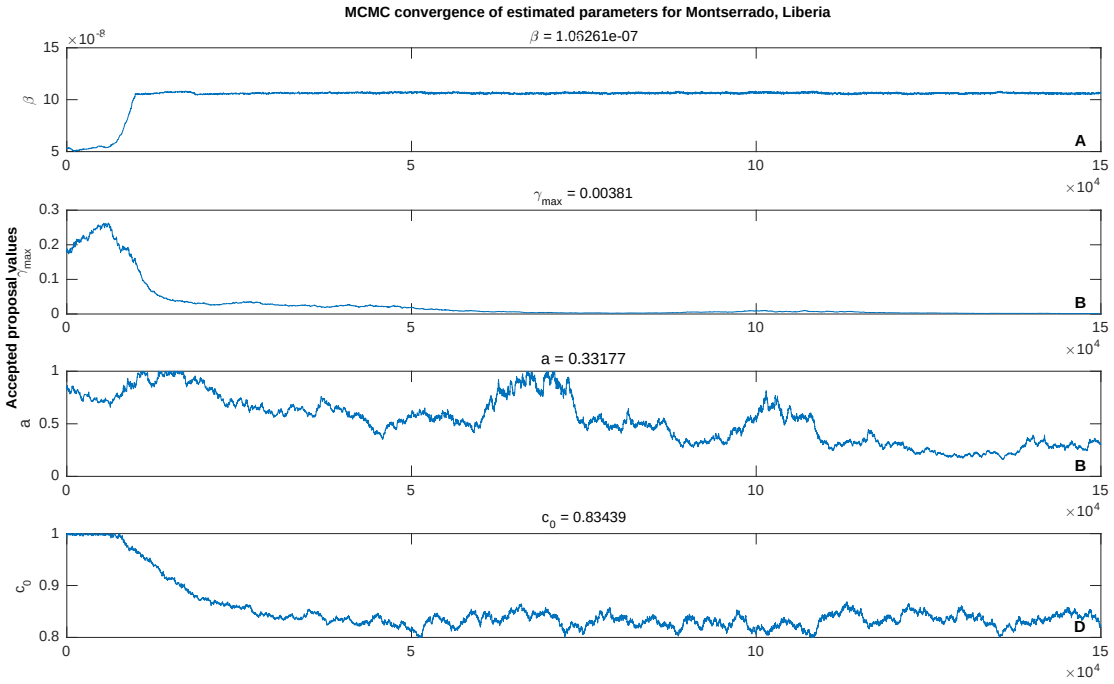

Figure C.40: MCMC convergence for Montserrado, Liberia.

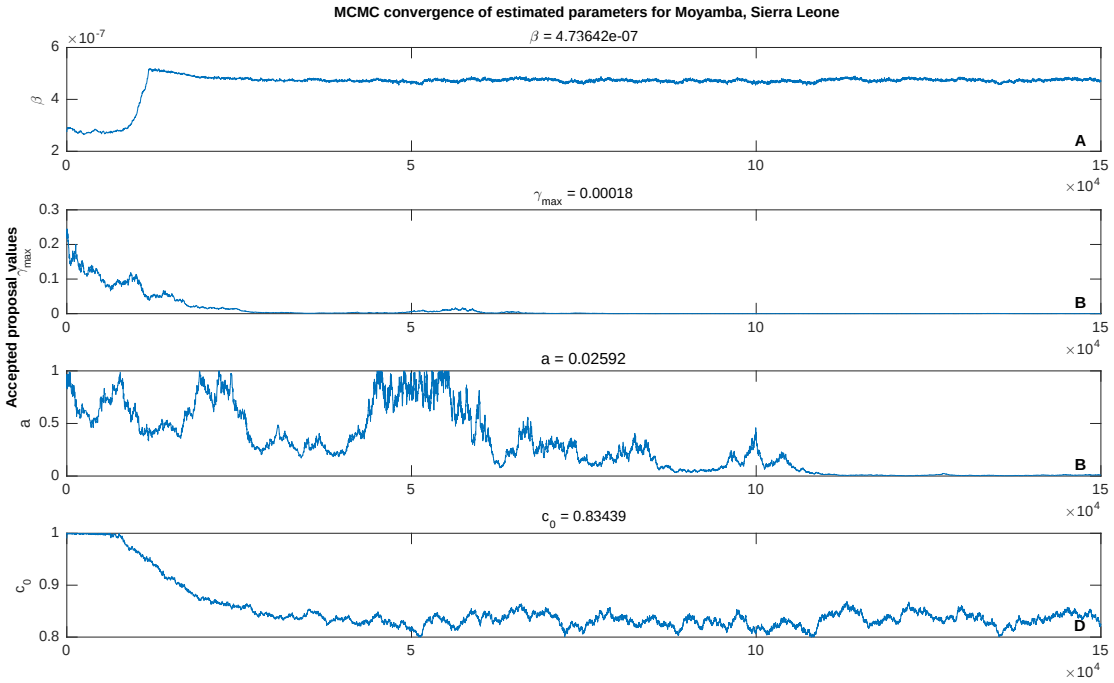

Figure C.41: MCMC convergence for Moyamba, Sierra Leone.

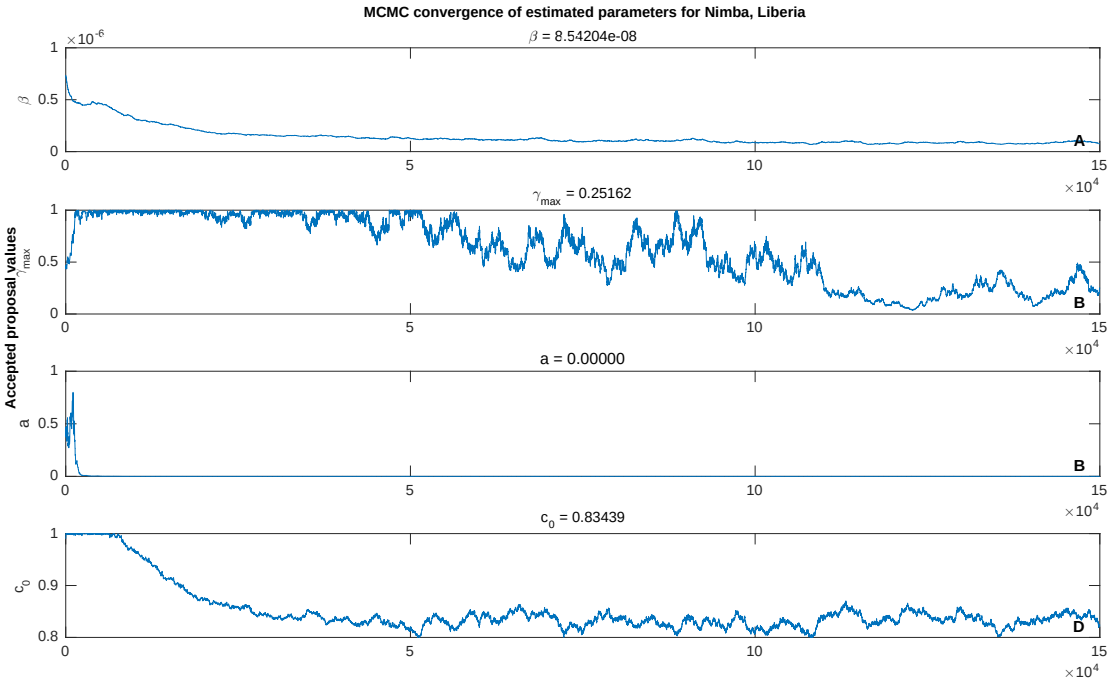

Figure C.42: MCMC convergence for Nimba, Liberia.

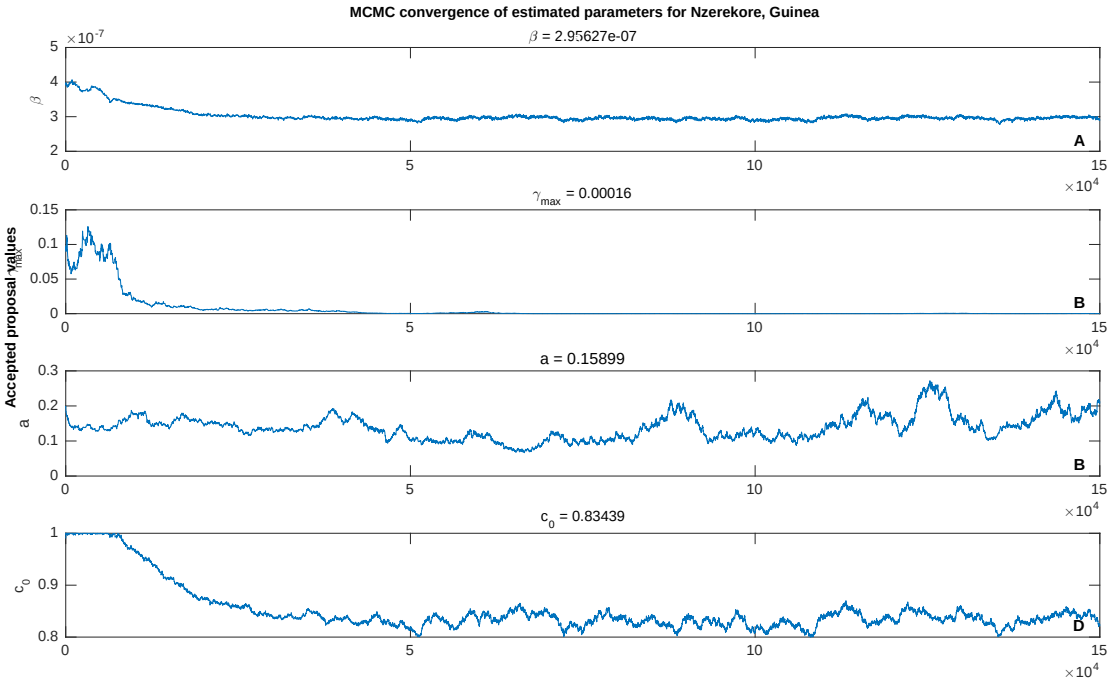

Figure C.43: MCMC convergence for Nzerekore, Guinea.

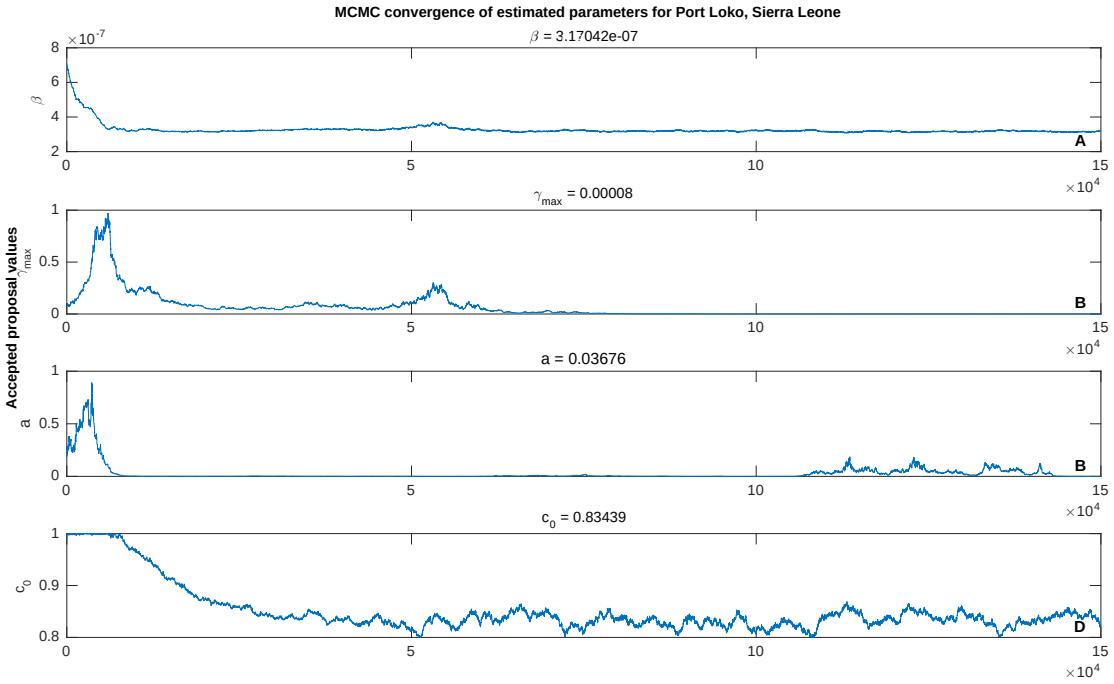

Figure C.44: MCMC convergence for Port Loko, Sierra Leone.

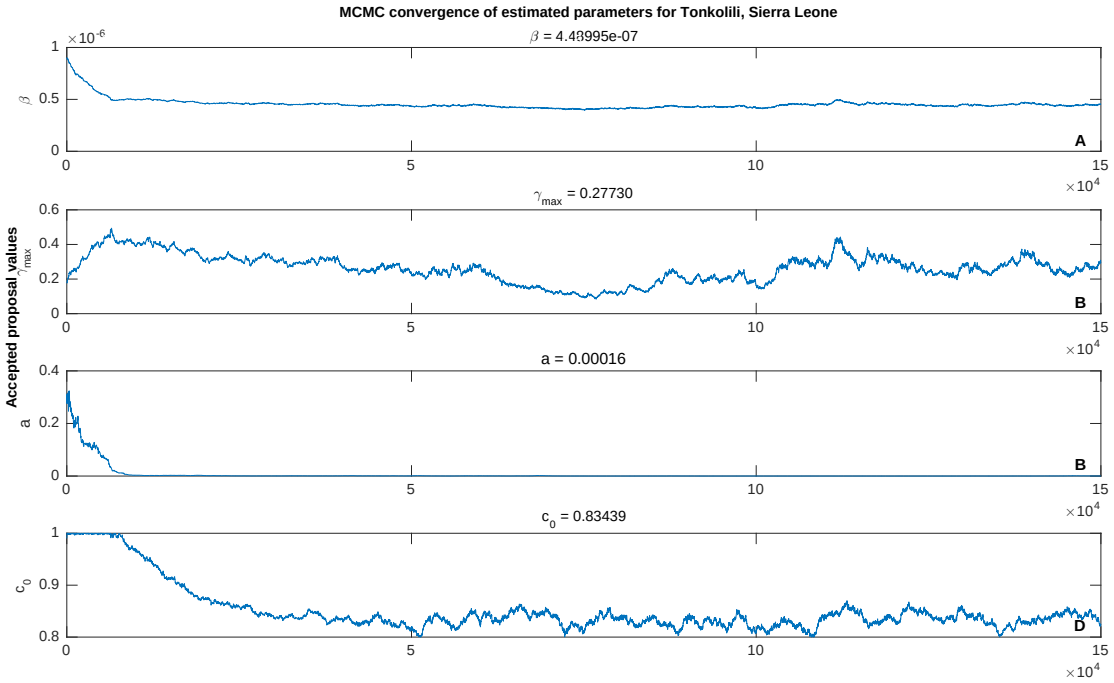

Figure C.45: MCMC convergence for Tonkolili, Sierra Leone.

Bibliography

ABDA. 2013. Die Apotheke Zahlen Daten Fakten 2013. URL http://www.abda.de/fileadmin/assets/ZDF/ZDF_2013/ABDA_ZDF_2013_Brosch.pdf. Retrieved on July 20th, 2014.

Abdollahian, Mehrnaz & Das, Tapas K. 2015. A MDP model for breast and ovarian cancer intervention strategies for BRCA1/2 mutation carriers. *IEEE journal of biomedical and health informatics* **19**(2) 720–727.

Alagoz, Oguzhan, Chhatwal, Jagpreet & Burnside, Elizabeth S. 2013. Optimal Policies for Reducing Unnecessary Follow-Up Mammography Exams in Breast Cancer Diagnosis. *Decision Analysis* **10**(3) 200–224.

Alagoz, Oguzhan, Hu, Hao, Schaefer, Andrew J. & Roberts, Mark S. 2010. Markov decision processes. *Medical Decision Making* **30**(4) 474–483.

Alagoz, Oguzhan, Maillart, Lisa M., Schaefer, Andrew J. & Roberts, Mark S. 2004. The Optimal Timing of Living-Donor Liver Transplantation. *Management Science* **50**(10) 1420–1430.

Alagoz, Oguzhan, Maillart, Lisa M., Schaefer, Andrew J. & Roberts, Mark S. 2007. Determining the Acceptance of Cadaveric Livers Using an Implicit Model of the Waiting List. *Operations Research* **55**(1) 24–36.

Alicke, Knut, Ebel, Thomas, Schrader, Ulf & Shah, Ketan, eds. 2014. *Finding Opportunity in Uncertainty: A New Paradigm for Pharmaceutical Supply Chains*. McKinsey & Company.

Althaus, Christian L. 2014. Estimating the reproduction number of Ebola virus during the 2014 outbreak in West Africa. *PLoS Currents: Outbreaks* **6**(1).

Anderson, Kristin, Jacobson, Judith S., Heitjan, Daniel F., Zivin, Joshua Graff, Hershman, Dawn, Neugut, Alfred I. & Grann, Victor R. 2006. Cost-Effectiveness of

Preventive Strategies for Women with a BRCA1 or a BRCA2 Mutation. *Annals of Internal Medicine* **144**(6) 397–406.

Anderson, Roy M. & May, Robert M. 1991. *Infectious Diseases of Humans*. Oxford University Press, Oxford and New York.

Angulo, Andres, Nachtmann, Heather & Waller, Matthew A. 2004. Supply Chain information sharing in a Vendor Managed Inventory Partnership. *Journal of Business Logistics* **25**(1) 101–120.

Antoniou, Antonis C., Pharoah, Paul D., Narod, Steven A. et al. 2003. Average risks of breast and ovarian cancer associated with brca1 or brca2 mutations detected in case series unselected for family history. *The American Journal of Human Genetics* **72**(5) 1117–1130.

Apolloni, Andrea, Poletto, Chiara, Ramasco, José J., Jensen, Pablo & Colizza, Vittoria. 2014. Metapopulation epidemic models with heterogeneous mixing and travel behaviour. *Theoretical Biology and Medical Modelling* **11**(3).

Aptel, Olivier & Pourjalali, Hamid. 2001. Improving activities and decreasing costs of logistics in hospitals: a comparison of US and French hospitals. *The International Journal of Accounting* **36**(1) 65–90.

Arino, Julien & Van Den Driessche, Pauline. 2003. The basic reproduction number in a multi-city compartmental epidemic model. *Positive Systems*, 135–142. Springer.

Armstrong, Katrina, Schwartz, J. Sanford, Randall, Thomas, Rubin, Stephen C. & Weber, Barbara. 2004. Hormone Replacement Therapy and Life Expectancy After Prophylactic Oophorectomy in Women With BRCA1/2 Mutations: A Decision Analysis. *Journal of Clinical Oncology* **22**(6) 1045–1054.

Aviv, Yossi. 2001. The effect of collaborative forecasting on supply chain performance. *Management Science* **47**(10) 1326–1343.

Aviv, Yossi. 2007. On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Science* **53**(5) 777–794.

Ayer, Turgay, Alagoz, Oguzhan & Stout, Natasha K. 2012. OR Forum—A POMDP Approach to Personalize Mammography Screening Decisions. *Operations Research* **60**(5) 1019–1034.

Balci, Osman. 1990. Guidelines for successful simulation studies. Osman. Balci, Randall P. Sadowski & Richard E. Nance, eds., *1990 Winter Simulation Conference proceedings*, 25–32.

Barlow, Richard E. & Proschan, Frank. 1965. *Mathematical Theory of Reliability*. John Wiley & Sons, New York.

Barton, Mary B., West, Carmen N., Liu, In-Lu A., Harris, Emily L., Rolnick, Sharon J., Elmore, Joann G., Herrinton, Lisa J., Greene, Sarah M., Nekhlyudov, Larissa, Fletcher, Suzanne W. & Geiger, Ann M. 2005. Complications Following Bilateral Prophylactic Mastectomy. *JNCI Monographs* **2005**(35) 61–66.

Barz, Christiane & Rajaram, Kumar. 2015. Elective patient admission and scheduling under multiple resource constraints. *Production and Operations Management* **24**(12) 1907–1930.

Bertsekas, Dimitri P. 1995. *Dynamic Programming and Optimal Control*, vol. 2. Athena Scientific Belmont, MA.

Bhakoo, Vikram, Singh, Prakash & Sohal, Amrik. 2012. Collaborative management of inventory in Australian hospital supply chains: practices and issues. *Supply Chain Management: An International Journal* **17**(2) 217–230.

Blount, Steven M., Galambosi, Agnes. & Yakowitz, Sidney J. 1997. Nonlinear and dynamic programming for epidemic intervention. *Applied Mathematics and Computation* **86**(2) 123–136.

Bolton, Kelly L., Chenevix-Trench, Georgia, Goh, Cindy, Sadetzki, Siegal, Ramus, Susan J., Karlan, Beth Y., Lambrechts, Diether, Despierre, Evelyn, Barrowdale, Daniel & McGuffog, Lesley. 2012. Association Between BRCA1 and BRCA2 Mutations and Survival in Women With Invasive Epithelial Ovarian Cancer. *JAMA* **307**(4) 382.

Brandeau, Margaret L. 2004. Allocating Resources to Control Infectious Diseases. Margaret L. Brandeau, François Sainfort & William P. Pierskalla, eds., *Operations Research and Health Care*, vol. 70 of *International Series in Operations Research & Management Science*, 443–464. Springer US.

Brandeau, Margaret L., Zaric, Gregory S. & Richter, Anke. 2003. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *Journal of Health Economics* **22**(4) 575–598.

Brown, Monica, Tsodikov, Alex, Bauer, Katrina R., Parise, Carol A. & Caggiano, Vincent. 2008. The role of human epidermal growth factor receptor 2 in the survival of women with estrogen and progesterone receptor–negative, invasive breast cancer: The California Cancer Registry, 1999–2004. *Cancer* **112**(4) 737–747.

Burns, Lawton R. 2002. *The health care value chain.* Jossey-Bass, San Francisco.

Butler, Declan. 2014a. Global Ebola response kicks into gear at last. *Nature* **513**(7519) 469.

Butler, Declan. 2014b. Models overestimate Ebola cases. *Nature* **515**(18) 7525.

Byrne, P. J. & Heavey, Cathal. 2006. The impact of information sharing and forecasting in capacitated industrial supply chains: A case study. *International Journal of Production Economics* **103**(1) 420–437.

Cachon, Gérard P. & Fisher, Marshall. 2000. Supply chain inventory management and the value of shared information. *Management Science* **46**(8) 1032–1048.

Cachon, Gérard P. & Lariviere, Martin A. 1999. Capacity choice and allocation: strategic behavior and supply chain performance. *Management Science* **45**(8) 1091–1108.

Centers for Disease Control and Prevention & National Center for Health Statistics. 2014. Underlying Cause of Death 1999-2014 on CDC WONDER Online Database, released 2015. URL http://wonder.cdc.gov/ucd-icd10.html. Retrieved on March 16$^{\text{th}}$, 2016.

Chen, Sining & Parmigiani, Giovanni. 2007. Meta-Analysis of BRCA1 and BRCA2 Penetrance. *Journal of Clinical Oncology* **25**(11) 1329–1333.

Chen, Tsung-Hui & Chen, Jen-Ming. 2005. Optimizing supply chain collaboration based on joint replenishment and channel coordination. *Transportation Research Part E: Logistics and Transportation Review* **41**(4) 261–285.

Chhatwal, Jagpreet, Alagoz, Oguzhan & Burnside, Elizabeth S. 2010. Optimal Breast Biopsy Decision-Making Based on Mammographic Features and Demographic Factors. *Operations Research* **58**(6) 1577–1591.

Chiang, W. K. & Feng, Y. 2007. The value of information sharing in the presence of supply uncertainty and demand volatility. *International Journal of Production Research* **45**(6) 1429–1447.

Chowell, Gerardo, Hengartner, Nick W., Castillo-Chavez, Carlos, Fenimore, Paul W. & Hyman, J. M. 2004. The basic reproductive number of Ebola and the effects of public health measures: the cases of Congo and Uganda. *Journal of Theoretical Biology* **229**(1) 119–126.

Currie, C. S.M. 2007. Bayesian methodology for dynamic modelling. *Journal of Simulation* **1**(2) 97–107.

Dacosta-Claro, Ivan. 2002. The performance of material management in health care organizations. *The International Journal of Health Planning and Management* **17**(1) 69–85.

Dancey, Janet E., Bedard, Philippe L., Onetto, Nicole & Hudson, Thomas J. 2012. The Genetic Basis for Cancer Treatment Decisions. *Cell* **148**(3) 409–420.

de Farias, Daniela Pucci & van Roy, Benjamin. 2003. The linear programming approach to approximate dynamic programming. *Operations Research* **51**(6) 850–865.

de Vries, Jan, Bhakoo, Vikram & Chan, Caroline. 2011. Collaborative implementation of e–business processes within the health–care supply chain. *Supply Chain Management: An International Journal* **16**(3) 184–193.

Del Valle, S., Hethcote, H., Hyman, J. M. & Castillo-Chavez, C. 2005. Effects of behavioral changes in a smallpox attack model. *Mathematical Biosciences* **195**(2) 228–251.

Doctors Without Borders. 2015. MSF Opens New Ebola Treatment Centers in Sierra Leone to Increase Access to Care. URL http://www.doctorswithoutborders.org/article/msf-opens-new-ebola-treatment-centers-sierra-leone-increase-access-care. Retrieved on August 13th, 2015.

Domchek, Susan M., Friebel, Tara M., Neuhausen, Susan L. et al. 2006. Mortality after bilateral salpingo-oophorectomy in BRCA1 and BRCA2 mutation carriers: a prospective cohort study. *The Lancet Oncology* **7**(3) 223–229.

Domchek, Susan M., Friebel, Tara M., Singer, Christian F. et al. 2010. Association of Risk-Reducing Surgery in BRCA1 or BRCA2 Mutation Carriers With Cancer Risk and Mortality. *JAMA* **304**(9) 967–975.

Duffy, S. W. & Nixon, R. M. 2002. Estimates of the likely prophylactic effect of tamoxifen in women with high risk BRCA1 and BRCA2 mutations. *British journal of cancer* **86**(2) 218–221.

Ebel, Thomas, George, Katy, Larsen, Erik & ungerman, Drew. 2013. Building New Strengths in the Healthcare Supply Chain. URL http://www.mckinsey.com/~/media/mckinsey/dotcom/client_service/Pharma%20and%20Medical%20Products/PMP%20NEW/PDFs/McKinsey%20white%20paper%20-%20building%20new%20strenghts%20in%20healthcare%20supply%20chain%20VF.ashx. Retrieved on March 4th, 2014.

Edge, Stephen B. & Compton, Carolyn C. 2010. The American Joint Committee on Cancer: the 7th Edition of the AJCC Cancer Staging Manual and the Future of TNM. *Annals of Surgical Oncology* **17**(6) 1471–1474.

Eisen, Andrea, Lubinski, Jan, Klijn, Jan, Moller, Pal, Lynch, Henry T., Offit, Kenneth, Weber, Barbara, Rebbeck, Tim, Neuhausen, Susan L. & Ghadirian, Parviz. 2005. Breast Cancer Risk Following Bilateral Oophorectomy in BRCA1 and BRCA2 Mutation Carriers. *Journal of Clinical Oncology* **23**(30) 7491–7496.

Fauci, A. S. & Morens, D. M. 2012. The perpetual challenge of infectious diseases. *The New England journal of medicine* **366**(5) 454–461.

Ferguson, Neil M., Cucunuba, Zulma M., Dorigatti, Ilaria, Nedjati-Gilani, Gemma L., Donnelly, Christl A., Basanez, Maria-Gloria, Nouvellet, Pierre & Lessler, Justin. 2016. Countering the zika epidemic in latin america. *Science* **353**(6297) 353–354.

Ferrari, Matthew J., Bjørnstad, Ottar N. & Dobson, Andrew P. 2005. Estimation and inference of R_0 of an infectious pathogen by a removal method. *Mathematical Biosciences* **198**(1) 14–26.

Fisman, David, Khoo, Edwin & Tuite, Ashleigh. 2014. Early epidemic dynamics of the West African 2014 Ebola outbreak: estimates derived with a simple two-parameter model. *PLoS Currents: Outbreaks* (6).

Fleisch, Elgar & Tellkamp, Christian. 2005. Inventory inaccuracy and supply chain performance. *International Journal of Production Economics* **95**(3) 373–385.

Fliedner, Gene. 2003. CPFR: an emerging supply chain tool. *Industrial management & Data systems* **103**(1) 14–21.

Frank, Thomas S., Deffenbaugh, Amie M., Reid, Julia E., Hulick, Mark, Ward, Brian E., Lingenfelter, Beth, Gumpper, Kathi L., Scholl, Thomas, Tavtigian, Sean V., Pruss, Dmitry R. & Critchfield, Gregory C. 2002. Clinical Characteristics of Individuals With Germline Mutations in BRCA1 and BRCA2: Analysis of 10,000 Individuals. *Journal of Clinical Oncology* **20**(6) 1480–1490.

Fry, Alison., Busby-Earle, Camille, Rush, Robert & Cull, Ann. 2001. Prophylactic oophorectomy versus screening: psychosocial outcomes in women at increased risk of ovarian cancer. *Psycho-oncology* **10**(3) 231–241.

Garcia, Christine, Wendt, Jacqueline, Lyon, Liisa, Jones, Jennifer, Littell, Ramey D., Armstrong, Mary Anne, Raine-Bennett, Tina & Powell, C. Bethan. 2014. Risk management options elected by women after testing positive for a BRCA mutation. *Gynecologic Oncology* **132**(2) 428–433.

Gass, Saul I. 1983. Feature Article—Decision-Aiding Models: Validation, Assessment, and Related Issues for Policy Analysis. *Operations Research* **31**(4) 603–631.

Geiger, Ann M., West, Carmen N., Nekhlyudov, Larissa, Herrinton, Lisa J., Liu, In-Liu A., Altschuler, Andrea, Rolnick, Sharon J., Harris, Emily L., Greene, Sarah M., Elmore, Joann G., Emmons, Karen M. & Fletcher, Suzanne W. 2006. Contentment with quality of life among breast cancer survivors with and without contralateral prophylactic mastectomy. *Journal of clinical oncology: official journal of the American Society of Clinical Oncology* **24**(9) 1350–1356.

Gerberry, David J., Wagner, Bradley G., Garcia-Lerma, J. Gerardo, Heneine, Walid & Blower, Sally. 2014. Using geospatial modelling to optimize the rollout of antiretroviral-based pre-exposure HIV interventions in Sub-Saharan Africa. *Nature communications* **5**(5454).

Gettleman, Jeffrey. 2014. As Ebola Rages, Poor Planning Thwarts Efforts. *The New York Times* URL http://www.nytimes.com/2014/12/07/world/africa/as-ebola-rages-in-sierra-leone-poor-planning-thwarts-efforts.html. Retrieved on May 5th, 2015.

Gibson, G. 1998. Estimating parameters in stochastic compartmental models using markov chain methods. *Mathematical Medicine and Biology* **15**(1) 19–40.

Gomes, Marcelo. F., Piontti, Ana. P., Rossi, Luca, Chao, Dennis, Longini, Ira & Halloran, M. Elizabeth. 2014. Assessing the international spreading risk associated with the 2014 west african ebola outbreak. *PLoS Currents: Outbreaks* .

Graeser, Monika K., Engel, Christoph, Rhiem, Kerstin et al. 2009. Contralateral breast cancer risk in BRCA1 and BRCA2 mutation carriers. *Journal of Clinical Oncology* **27**(35) 5887–5892.

Grann, Victor R., Jacobson, Judith S., Sundararajan, Vijaya, Albert, Steven M., Troxel, Andreas B. & Neugut, Alfred I. 1999a. The quality of life associated with prophylactic treatments for women with BRCA1/2 mutations. *The cancer journal from Scientific American* **5**(5) 283–292.

Grann, Victor R., Jacobson, Judith S., Thomason, D., Hershman, Dawn, Heitjan, Daniel F. & Neugut, A. I. 2002. Effect of Prevention Strategies on Survival and Quality-Adjusted Survival of Women With BRCA1/2 Mutations. *Journal of Clinical Oncology* **20**(10) 2520–2529.

Grann, Victor R., Panageas, Katherine S., Whang, William, Antman, Karen H. & Neugut, Alfred I. 1998. Decision analysis of prophylactic mastectomy and oophorectomy in BRCA1-positive or BRCA2-positive patients. *Journal of Clinical Oncology* **16**(3) 979–985.

Grann, Victor R., Patel, Priya, Bharthuar, Anubha, Jacobson, Judith S., Warner, Ellen, Anderson, Kristin, Warner, Eiran, Tsai, Wei-Yann, Hill, Kimberly A., Neugut, Alfred I. & Hershman, Dawn. 2010. Breast cancer-related preferences among women with and without BRCA mutations. *Breast cancer research and treatment* **119**(1) 177–184.

Grann, Victor R., Whang, William, Jacobson, Judith S., Heitjan, Daniel F., Antman, Karen H. & Neugut, Alfred I. 1999b. Benefits and costs of screening Ashkenazi Jewish women for BRCA1 and BRCA2. *Journal of Clinical Oncology* **17**(2) 494.

Grant, Robert L. 2014. Converting an odds ratio to a range of plausible relative risks for better communication of research findings. *BMJ* **348** f7450.

Greenhalgh, David. 1986. Control of an epidemic spreading in a heterogeneously mixing population. *Mathematical Biosciences* **80**(1) 23–45.

Gronwald, J., Robidoux, A., Kim-Sing, C. et al. 2014. Duration of tamoxifen use and the risk of contralateral breast cancer in BRCA1 and BRCA2 mutation carriers. *Breast cancer research and treatment* **146**(2) 421–427.

Groseth, Allison, Feldmann, Heinz & Strong, James E. 2007. The ecology of Ebola virus. *Trends Microbiology* **15**(9) 408–416.

Haavik, Stan. 2000. Building a demand-driven, vendor-managed supply chain. *Healthcare financial management: journal of the Healthcare Financial Management Association* **54**(2) 56–61.

Hair, Joseph F., Tatham, Ronald L., Anderson, Rolph E. & Black, William. 2006. *Multivariate data analysis*. Pearson Prentice Hall Upper Saddle River, NJ, 7 edn.

Hall, Jeff M., Lee, Ming K., Newman, Beth, Morrow, Jan E., Anderson, Lee A., Huey, Bing & King, Mary-Claire. 1990. Linkage of Early-Onset Familial Breast Cancer to Chromosome 17q21. *Science* **250**(4988) 1684–1689.

Hall, Michael J., Reid, Julia E., Burbidge, Lynn A., Pruss, Dmitry, Deffenbaugh, Amie M., Frye, Cynthia, Wenstrup, Richard J., Ward, Brian E., Scholl, Thomas A. & Noll, Walter W. 2009. BRCA1 and BRCA2 mutations in women of different ethnicities undergoing testing for hereditary breast–ovarian cancer. *Cancer* **115**(10) 2222–2233.

Harmsen, Marline G., Hermens, Rosella P.M.G., Prins, Judith B., Hoogerbrugge, Nicoline & de Hullu, Joanne A. 2015. How medical choices influence quality of life of women carrying a BRCA mutation. *Critical Reviews in Oncology/Hematology* **96**(3) 555–568.

Heesterbeek, Hans, Anderson, Roy M., Andreasen, Viggo et al. 2015. Modeling infectious disease dynamics in the complex landscape of global health. *Science* **347**(6227) aaa4339.

Henschke, C., Sundmacher, L. & Busse, R. 2013. Structural changes in the German pharmaceutical market: price setting mechanisms based on the early benefit evaluation. *Health policy (Amsterdam, Netherlands)* **109**(3) 263–269.

Hochbaum, Dorit S. 1995. A nonlinear knapsack problem. *Operations Research Letters* **17**(3) 103–110.

Holt, Charles C. 2004. Forecasting seasonals and trends by exponentially weighted moving averages. *International journal of forecasting* **20**(1) 5–10.

Howlader, N., Noone, A. M., Krapcho, M., Neyman, N., Aminou, R., Waldron, W., Altekruse, S. F., Kosary, C. L., Ruhl, J., Tatalovich, Z., Cho, H., Mariotto, A., Eisner, M. P., Lewis, D. R., Chen, H. S., Feuer, E. J. & Cronin, K. A. 2012. *SEER Cancer Statistics Review, 1975-2009 (Vintage 2009 Populations)*. National Cancer Institute, Bethesda, MD. URL http://seer.cancer.gov/csr/1975_2009_pops09/. Retrieved on June 27th, 2015.

Hulshof, PeterJ.H., Mes, MartijnR.K., Boucherie, RichardJ. & Hans, ErwinW. 2015. Patient admission planning using Approximate Dynamic Programming. *Flexible Services and Manufacturing Journal* 1–32.

Humanitarian Data Exchange. 2015a. Number of Ebola Cases and Deaths in Affected Countries. URL https://data.hdx.rwlabs.org/dataset/ebola-cases-2014. Retrieved on February 2nd,2015.

Humanitarian Data Exchange. 2015b. Sub-national time series data on Ebola cases and deaths in Guinea, Liberia, Sierra Leone, Nigeria, Senegal and Mali since March 2014. URL https://data.hdx.rwlabs.org/dataset/rowca-ebola-cases. Retrieved on January 22nd, 2015.

Inderfurth, Karl & Minner, Stefan. 1998. Safety stocks in multi-stage inventory systems under different service measures. *European Journal of Operational Research* **106**(1) 57–73.

Institut National de la Statistique. 2014. Répartition population et densité par Région et Préfecture. URL http://www.stat-guinee.org/index.php/statistiques/donnees-structurelles/demographie/31-population-densite-region-prefecture. Retrieved on February 2nd, 2015.

Jarrett, P. Gary. 1998. Logistics in the health care industry. *International Journal of Physical Distribution & Logistics Management* **28**(9/10) 741–772.

Kamalapur, Raj, Lyth, David & Houshyar, Azim. 2013. Benefits of CPFR and VMI Collaboration Strategies: a Simulation Study. *Journal of Operations and Supply Chain Management* **6**(2) 59–73.

Kanchanasuntorn, Kanchana & Techanitisawad, Anulark. 2006. An approximate periodic model for fixed-life perishable products in a two-echelon inventory–distribution system. *International Journal of Production Economics* **100**(1) 101–115.

Kanda, Arun & Deshmukh, S. G. 2008. Supply chain coordination: perspectives, empirical studies and research directions. *International Journal of Production Economics* **115**(2) 316–335.

Karaesmen, Itir Z., Scheller-Wolf, Alan & Deniz, Borga. 2011. Managing Perishable and Aging Inventories: Review and Future Research Directions. Karl G. Kempf, Pınar Keskinocak & Reha Uzsoy, eds., *Planning Production and Inventories in the Extended Enterprise*, vol. 151 of *International Series in Operations Research & Management Science*, 393–436. Springer US, Boston, MA.

Katsanis, Sara Huston & Katsanis, Nicholas. 2013. Molecular genetic testing and the future of clinical genomics. *Nature reviews. Genetics* **14**(6) 415–426.

Kauff, Noah D., Domchek, Susan M., Friebel, Tara M. et al. 2008. Risk-Reducing Salpingo-Oophorectomy for the Prevention of BRCA1- and BRCA2-Associated Breast and Gynecologic Cancer: A Multicenter, Prospective Study. *Journal of Clinical Oncology* **26**(8) 1331–1337.

Kellerer, Hans, Pferschy, Ulrich & Pisinger, D. 2004. *Knapsack Problems*. Springer, Berlin and New York.

Ketzenberg, Michael & Ferguson, Mark E. 2008. Managing Slow-Moving Perishables in the Grocery Industry. *Production and Operations Management* **17**(5) 513–521.

Khademi, Amin, Saure, Denis R., Schaefer, Andrew J., Braithwaite, Ronald S. & Roberts, Mark S. 2015. The Price of Nonabandonment. *Manufacturing & Service Operations Management* **17**(4) 554–570.

Khan, Adnan, Naveed, Mahim, Dur-e Ahmad, Muhammad & Imran, Mudassar. 2015. Estimating the basic reproductive ratio for the Ebola outbreak in Liberia and Sierra Leone. *Infectious Diseases of Poverty* **4**(1) 13.

Kiely, Dan. 2004. The state of pharmaceutical industry supply planning and demand Forecasting. *Journal of Business Forecasting Methods & Systems* **23**(3).

King, M.-C. 2003. Breast and Ovarian Cancer Risks Due to Inherited Mutations in BRCA1 and BRCA2. *Science* **302**(5645) 643–646.

King, Mary-Claire, Wieand, Sam, Hale, Kathryn, Lee, Ming, Walsh, Tom, Owens, Kelly, Tait, Jonathan, Ford, Leslie, Dunn, Barbara K., Costantino, Joseph, Wickerham, Lawrence, Wolmark, Norman & Fisher, Bernard. 2001. Tamoxifen and Breast Cancer Incidence Among Women With Inherited Mutations in BRCA1 and BRCA2: National Surgical Adjuvant Breast and Bowel Project (NSABP-P1) Breast Cancer Prevention Trial. *JAMA* **286**(18) 2251–2256.

Kurian, Allison W., Lichtensztajn, Daphne Y., Keegan, Theresa H. M., Nelson, David O., Clarke, Christina A. & Gomez, Scarlett L. 2014. Use of and mortality after bilateral mastectomy compared with other surgical treatments for breast cancer in California, 1998-2011. *JAMA* **312**(9) 902–914.

Kurian, Allison W., Munoz, Diego F., Rust, Peter, Schackmann, Elizabeth A., Smith, Michael, Clarke, Lauren, Mills, Meredith A. & Plevritis, Sylvia K. 2012. Online tool to guide decisions for BRCA1/2 mutation carriers. *Journal of Clinical Oncology* **30**(5) 497–506.

Lau, Jason S. K., Huang, George Q. & Mak, K. L. 2004. Impact of information sharing on inventory replenishment in divergent supply chains. *International Journal of Production Research* **42**(5) 919–941.

Lee, Larissa J., Alexander, Brian, Schnitt, Stuart J., Comander, Amy, Gallagher, Bridget, Garber, Judy E. & Tung, Nadine. 2011. Clinical outcome of triple negative breast cancer in BRCA1 mutation carriers and noncarriers. *Cancer* **117**(14) 3093–3100.

Legrand, Judith, Grais, Rebecca F., Boelle, Pierre Y., Valleron, Alain J. & Flahault, Antione. 2007. Understanding the dynamics of ebola epidemics. *Epidemiology and infection* **135**(04) 610–621.

Lekone, Phenyo E. & Finkenstädt, Bärbel F. 2006. Statistical inference in a stochastic epidemic SEIR model with control intervention: Ebola as a case study. *Biometrics* **62**(4) 1170–1177.

Lewnard, Joseph A., Ndeffo Mbah, Martial L., Alfaro-Murillo, Jorge A., Altice, Frederick L., Bawo, Luke, Nyenswah, Tolbert G. & Galvani, Alison P. 2014. Dynamics and control of Ebola virus transmission in Montserrado, Liberia: a mathematical modelling analysis. *The Lancet Infectious Diseases* **14**(12) 1189–1195.

Li, Jingquan, Shaw, Michael J., Sikora, Riyaz T., Tan, Gek Woo & Yang, Rachel. 2001. The effects of information sharing strategies on supply chain performance. *IEEE Transactions of Engineering Management* .

Liberia Institute of Statistics and Geo-Information Service. 2008. Liberia - General Information. URL http://www.geohive.com/cntry/liberia.aspx. Retrieved on February 2nd, 2015.

Lightfoot, Nigel, Rweyemamu, Mark & Heymann, David L. 2013. Preparing for the next pandemic. *BMJ* **346** f364.

Lobbezoo, D. J. A., van Kampen, R J W, Voogd, A. C., Dercksen, M. W., van den Berkmortel, F., Smilde, T. J., van de Wouw, A J, Peters, F. P. J., van Riel, J M G H, Peters, N A J B, de Boer, M., Peer, P. G. M. & Tjan-Heijnen, V. C. G. 2015. Prognosis of metastatic breast cancer: are there differences between patients with de novo and recurrent metastatic breast cancer? *British journal of cancer* **112**(9) 1445–1451.

Loftus, Peter. 2016. Drugmakers Raise Prices Despite Criticisms. *Wall Street Journal* URL http://www.wsj.com/articles/drugmakers-raise-prices-despite-criticisms-1452474210. Retrieved on April 3rd, 2016.

Long, Elisa F. & Ganz, Patricia A. 2015. Cost-effectiveness of Universal BRCA1/2 Screening: Evidence-Based Decision Making. *JAMA Oncology* **1**(9) 1217–1218.

Madalinska, Joanna B., Hollenstein, Judith, Bleiker, Eveline, van Beurden, Marc, Valdimarsdottir, Heiddis B., Massuger, Leon F., Gaarenstroom, Katja N., Mourits, Marian J., Verheijen, René H.M., van Dorst, Eleonora B., van der Putten, Hans, van der Velden, Ko, Boonstra, Henk & Aaronson, Neil K. 2005. Quality-of-life

effects of prophylactic salpingo-oophorectomy versus gynecologic screening among women at increased risk of hereditary ovarian cancer. *Journal of clinical oncology* **23**(28) 6890–6898.

Malone, Kathleen. E., Daling, Janet R., Doody, David R. et al. 2006. Prevalence and predictors of brca1 and brca2 mutations in a population-based study of breast cancer in white and black american women ages 35 to 64 years. *Cancer research* **66**(16) 8297–8308.

Mangili, Alexandra & Gendreau, Mark A. 2012. Transmission of infectious diseases during commercial air travel. *The Lancet* **365**(9463) 989–996.

Mason, Jennifer E., Denton, Brian T., Shah, Nilay D. & Smith, Steven A. 2014. Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *European Journal of Operational Research* **233**(3) 727–738.

Mason, Jennifer E., England, Darin A., Denton, Brian T., Smith, Steven A., Kurt, Murat & Shah, Nilay D. 2011. Optimizing Statin Treatment Decisions for Diabetes Patients in the Presence of Uncertain Future Adherence. *Medical Decision Making* **32**(1) 154–166.

Mavaddat, Nasim, Barrowdale, Daniel, Andrulis, Irene L. et al. 2012. Pathology of breast and ovarian cancers among BRCA1 and BRCA2 mutation carriers: results from the Consortium of Investigators of Modifiers of BRCA1/2 (CIMBA). *Cancer epidemiology, biomarkers & prevention : a publication of the American Association for Cancer Research, cosponsored by the American Society of Preventive Oncology* **21**(1) 134–147.

McCarthy, Teresa M. & Golicic, Susan L. 2002. Implementing collaborative forecasting to improve supply chain performance. *International Journal of Physical Distribution & Logistics Management* **32**(6) 431–454.

McDonnell, Shannon K., Schaid, Daniel J., Myers, Jeffrey L., Grant, Clive S., Donohue, John H., Woods, John E., Frost, Marlene H., Johnson, Joanne L., Sitta, Diana L., Slezak, Jeffrey M., Crotty, Thomas B., Jenkins, Robert B., Sellers, Thomas A. & Hartmann, Lynn C. 2001. Efficacy of Contralateral Prophylactic Mastectomy in Women With a Personal and Family History of Breast Cancer. *Journal of Clinical Oncology* **19**(19) 3938–3943.

McKone-Sweet, Kathleen E., Hamilton, Paul & Willis, Susan B. 2005. The ailing healthcare supply chain: a prescription for change. *Journal of Supply Chain Management* **41**(1) 4–17.

Medlock, Jan & Galvani, Alison P. 2009. Optimizing influenza vaccine distribution. *Science* **325**(5948) 1705–1708.

Meltzer, Martin I., Atkins, Charisma Y., Santibanez, Scott, Knust, Barbara, Petersen, Brett W., Ervin, Elizabeth D., Nichol, Stuart T., Damon, Inger K. & Washington, Michael L. 2014. Estimating the future number of cases in the Ebola epidemic — Liberia and Sierra Leone, 2014–2015. *MMWR* **63**(3) 1–14.

Merler, Stefano, Ajelli, Marco, Fumanelli, Laura, Gomes, Marcelo F. C., Piontti, Ana Pastore, Rossi, Luca, Chao, Dennis L., Longini, Ira M., Halloran, M. Elizabeth & Vespignani, Alessandro. 2015. Spatiotemporal spread of the 2014 outbreak of Ebola virus disease in Liberia and the effectiveness of non-pharmaceutical interventions: a computational modelling analysis. *The Lancet Infectious Diseases* **15**(2) 204–211.

Michelsen, Trond M., Dørum, Anne, Tropé, Claes G., Fosså, Sophie D. & Dahl, Alv A. 2009. Fatigue and quality of life after risk-reducing salpingo-oophorectomy in women at increased risk for hereditary breast-ovarian cancer. *International journal of gynecological cancer : official journal of the International Gynecological Cancer Society* **19**(6) 1029–1036.

Minner, Stefan. 1997. Dynamic programming algorithms for multi-stage safety stock optimization. *OR Spektrum* **19**(4) 261–271.

Morens, David M., Folkers, Gregory K. & Fauci, Anthony S. 2004. The challenge of emerging and re-emerging infectious diseases. *Nature* **430**(6996) 242–249.

Moyer, Virginia A. 2014. Risk Assessment, Genetic Counseling, and Genetic Testing for BRCA-Related Cancer in Women: U.S. Preventive Services Task Force Recommendation Statement. *Annals of Internal Medicine* **160**(4) 271–281.

National Comprehensive Cancer Network. 2016. NCCN Clinical Practice Guidelines in Concology. Genetic/Familial Assessment: Breast and Ovarian. URL https://www.nccn.org/professionals/physician_gls/pdf/genetics_screening.pdf. Retrieved on July 24$^{\text{th}}$, 2016.

Ndeffo Mbah, Martial L. & Gilligan, Christopher A. 2011. Resource allocation for epidemic control in metapopulations. *PloS one* **6**(9).

Nilsson, Martin P., Hartman, Linda, Kristoffersson, Ulf, Johannsson, Oskar T., Borg, Åke, Henriksson, Karin, Lanke, Elsa, Olsson, Håkan & Loman, Niklas. 2014. High risk of in-breast tumor recurrence after BRCA1/2-associated breast cancer. *Breast cancer research and treatment* **147**(3) 571–578.

Nohdurft, Eike, Long, Elisa F. & Spinler, Stefan. 2015a. Efficient Spatial Allocation of Epidemic Intervention Resources with a Focus on Ebola in West Africa. *INFORMS Annual Meeting (November 2).* Philadelphia, PA, USA.

Nohdurft, Eike, Long, Elisa F. & Spinler, Stefan. 2015b. Optimal Allocation of Epidemic Intervention Resources with an Application to the 2014 West Africa Ebola Crisis. *POMS Annual Conference (May 9).* Washington, DC, USA.

Nohdurft, Eike, Long, Elisa F. & Spinler, Stefan. 2016a. Spatial allocation of resources for epidemic control: A greedy heuristic versus dynamic programming approach. Unpublished manuscript.

Nohdurft, Eike, Long, Elisa F. & Spinler, Stefan. 2016b. Was Angelina right? Optimizing cancer prevention strategies for BRCA carriers. Unpublished manuscript.

Nohdurft, Eike & Spinler, Stefan. 2014a. Coordinating the Health Care Supply Chain - Saving costs and improving service level through joint Forecasting and Planning. *International conference on Operations Research.* Aachen, Germany.

Nohdurft, Eike & Spinler, Stefan. 2014b. Health Care Supply Chain Coordination – Service Improvement and Cost Saving through Collaboration. *INFORMS Annual Meeting (November 12).* San Francisco, CA, USA.

Nohdurft, Eike & Spinler, Stefan. 2016. Evaluating benefits of coordination schemes in pharmaceutical supply chains. Unpublished manuscript under review at Decision Sciences.

OECD. 2016. Pharmaceutical spending (indicator). URL https://data.oecd.org/healthres/pharmaceutical-spending.htm. Retrieved on June 10th, 2016.

Parise, Carol A. & Caggiano, Vincent. 2014. Breast Cancer Survival Defined by the ER/PR/HER2 Subtypes and a Surrogate Classification according to Tumor Grade

and Immunohistochemical Biomarkers. *Journal of Cancer Epidemiology* **2014**(2) 1–11.

PHAGRO. 2016. Zahlen, Daten, Fakten. URL http://www.phagro.de/pharma-grosshandel/zahlen-daten-fakten/. Retrieved on February 2nd, 2016.

Philips, Mit & Markham, áine. 2014. Ebola: a failure of international collective action. *The Lancet* **384**(9949) 1181.

Portnoy, D. B., Roter, D. & Erby, L. H. 2010. The role of numeracy on client knowledge in BRCA genetic counseling. *Patient education and counseling* **81**(1) 131–136.

Powell, Warren B. 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, 2nd edn.

Puterman, Martin L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Rebbeck, Timothy R., Friebel, Tara, Lynch, Henry T., Neuhausen, Susan L., van 't Veer, Laura, Garber, Judy E., Evans, Gareth R., Narod, Steven A., Isaacs, Claudine, Matloff, Ellen, Daly, Mary B., Olopade, Olufunmilayo I. & Weber, Barbara L. 2004. Bilateral prophylactic mastectomy reduces breast cancer risk in BRCA1 and BRCA2 mutation carriers: the PROSE Study Group. *Journal of Clinical Oncology* **22**(6) 1055–1062.

Rebbeck, Timothy R., Kauff, Noah D. & Domchek, Susan M. 2009. Meta-analysis of Risk Reduction Estimates Associated With Risk-Reducing Salpingo-oophorectomy in BRCA1 or BRCA2 Mutation Carriers. *Journal of the National Cancer Institute* **101**(2) 80–87.

Reyna, Valerie F., Nelson, Wendy L., Han, Paul K. & Pignone, Michael P. 2015. Decision making and cancer. *The American psychologist* **70**(2) 105–118.

Richter, Anke, Brandeau, Margaret L. & Owens, Douglas K. 1999. An Analysis of Optimal Resource Allocation for Prevention of Infection with Human Immunodeficiency Virus (HIV) in Injection Drug Users and Non-Users. *Medical Decision Making* **19**(2) 167–179.

Rivard-Royer, Hugo, Landry, Sylvain & Beaulieu, Martin. 2002. Hybrid stockless: A case study: Lessons for health-care supply chain integration. *International Journal of Operations & Production Management* **22**(4) 412–424.

Rivers, Caitlin M., Lofgren, Eric T., Marathe, Madhav, Eubank, Stephen & Lewis, Bryan L. 2014. Modeling the impact of interventions on an epidemic of Ebola in Sierra Leone and Liberia. *PLoS Currents: Outbreaks* (2).

Robson, Mark E., Chappuis, Pierre O., Satagopan, Jaya, Wong, Nora, Boyd, Jeff, Goffin, John R., Hudis, Clifford, Roberge, David, Norton, Larry, Bégin, Louis R., Offit, Kenneth & Foulkes, William D. 2003. A combined analysis of outcome following breast cancer: differences in survival based on BRCA1/BRCA2 mutation status and administration of adjuvant treatment. *Breast Cancer Res* **6**(1) 1–10.

Romero, Alejandro. 2013. Managing Medicines in the Hospital Pharmacy: Logistics Inefficiencies. S.i. Ao, Craig Douglas, W. S. Grundfest & Jon Burgstone, eds., *Proceedings of the World Congress on Engineering and Computer Science 2013*, 1120–1125.

Rowthorn, Robert E., Laxminarayan, Ramanan & Gilligan, Christopher A. 2009. Optimal control of epidemics in metapopulations. *Journal of the Royal Society, Interface* **6**(41) 1135–1144.

Sandıkçı, Burhaneddin, Maillart, Lisa M., Schaefer, Andrew J., Alagoz, Oguzhan & Roberts, Mark S. 2008. Estimating the Patient's Price of Privacy in Liver Transplantation. *Operations Research* **56**(6) 1393–1410.

Sargent, G., Robert. 2013. Verification and validation of simulation models. *Journal of Simulation* **7**(1) 12–24.

Sari, Kazim. 2008. On the benefits of CPFR and VMI: A comparative simulation study. *International Journal of Production Economics* **113**(2) 575–586.

Sawik, Tadeusz. 2009. Coordinated supply chain scheduling. *International Journal of Production Economics* **120**(2) 437–451.

Scheel, Oliver, Eitelwein, Oliver & Moliner, Pablo. 2014. Preparing the Supply Chain Pharma Needs. URL https://www.atkearney.de/documents/856314/4686406/Preparing+the+Supply+Chain+Pharma+Needs.pdf/ef32535d-c6f6-4227-b5f7-283b5611d74d. Retrieved on May 10th, 2016.

Schneller, Eugene Stewart & Smeltzer, Larry R. 2006. *Strategic management of the health care supply chain.* Jossey-Bass, San Francisco, CA, 1st ed edn.

Schrag, Deborah. 2000. Life Expectancy Gains From Cancer Prevention Strategies for Women With Breast Cancer and BRCA1 or BRCA2 Mutations. *JAMA* **283**(5) 617.

Schrag, Deborah, Kuntz, Karen M., Garber, Judy E. & Weeks, Jane C. 1997. Decision Analysis — Effects of Prophylactic Mastectomy and Oophorectomy on Life Expectancy among Women with BRCA1 or BRCA2 Mutations. *New England Journal of Medicine* **336**(20) 1465–1471.

Shechter, Steven M., Bailey, Matthew D., Schaefer, Andrew J. & Roberts, Mark S. 2008. The Optimal Time to Initiate HIV Therapy Under Ordered Health States. *Operations Research* **56**(1) 20–33.

Shin, Hojung & Benton, W. C. 2004. Quantity Discount-Based Inventory Coordination. *Production and Operations Management* **13**(1) 63–76.

Shuster, Lynne T., Gostout, Bobbie S., Grossardt, Brandon R. & Rocca, Walter A. 2008. Prophylactic oophorectomy in premenopausal women and long-term health. *Menopause international* **14**(3) 111–116.

Siddiqui, M. & Rajkumar, S. V. 2012. The High Cost of Cancer Drugs and What We Can Do About It. *Mayo Clinic Proceedings* **87**(10) 935–943.

Sierra Leone Statistics. 2004. Final Results 2004 Population and Housing Census. URL http://www.sierra-leone.org/Census/ssl_final_results.pdf. Retrieved on February 2nd, 2015.

Silver, Edward A. & Peterson, Rein. 1985. *Decision systems for inventory management and production planning*. Wiley series in production/operations management. Wiley, New York, 2nd ed edn.

Simpson, K. F. 1958. In-process inventories. *Operations Research* **6**(6) 863–873.

Skjoett-Larsen, Tage, Thernøe, Christian & Andresen, Claus. 2003. Supply chain collaboration: theoretical perspectives and empirical evidence. *International Journal of Physical Distribution & Logistics Management* **33**(6) 531–549.

Sørlie, Therese, Tibshirani, Robert, Parker, Joel, Hastie, Trevor, Marron, J. S., Nobel, Andrew, Deng, Shibing, Johnsen, Hilde, Pesich, Robert, Geisler, Stephanie, Demeter, Janos, Perou, Charles M., Lønning, Per E., Brown, Patrick O., Børresen-Dale,

Anne-Lise & Botstein, David. 2003. Repeated observation of breast tumor subtypes in independent gene expression data sets. *Proceedings of the National Academy of Sciences of the United States of America* **100**(14) 8418–8423.

Tako, Antuela A. & Robinson, Stewart. 2012. The application of discrete event simulation and system dynamics in the logistics and supply chain context. *Decision Support Systems* **52**(4) 802–815.

Tercyak, Kenneth P., Peshkin, Beth N., Brogan, Barbara M., DeMarco, Tiffani A., Pennanen, Marie F., Willey, Shawna C., Magnant, Colette M., Rogers, Sarah, Isaacs, Claudine & Schwartz, Marc D. 2007. Quality of life after contralateral prophylactic mastectomy in newly diagnosed high-risk breast cancer patients who underwent BRCA1/2 gene testing. *Journal of clinical oncology : official journal of the American Society of Clinical Oncology* **25**(3) 285–291.

Terwiesch, Christian, Ren, Z. Justin, Ho, Teck H. & Cohen, Morris A. 2005. An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Science* **51**(2) 208–220.

Thomas, Andrea. 2016. Germany Mulls Limiting Prices Drug Firms Can Charge to Health System. URL http://www.wsj.com/articles/germany-mulls-limiting-prices-drug-firms-can-charge-to-health-system-1461307437. Retrieved on May 10th, 2016.

Trainer, Alison H., Lewis, Craig R., Tucker, Kathy, Meiser, Bettina, Friedlander, Michael & Ward, Robyn L. 2010. The role of BRCA mutation testing in determining breast cancer therapy. *Nature Reviews Clinical Oncology* **7**(12) 708–717.

Tsay, Andy A. 1999. The quantity flexibility contract and supplier-customer incentives. *Management Science* **45**(10) 1339–1358.

United Nations. 2015. Funding the Ebola Response. URL http://ebolaresponse.un.org/funding-ebola-response. Retrieved on May 14th, 2016.

UNMEER. 2015. Ebola Treatment Centers or Units (ETCs or ETUs). URL https://data.hdx.rwlabs.org/dataset/ebola-treatment-centers. Retrieved on March, 12th, 2015.

van den Driessche, P. & Watmough, James. 2002. Reproduction numbers and subthreshold endemic equilibria for compartmental models of disease transmission. *Mathematical Biosciences* **180**(1–2) 29–48.

van Gorp, T., Cadron, I., Despierre, E., Daemen, A., Leunen, K., Amant, F., Timmerman, D., Moor, B. De & Vergote, I. 2011. HE4 and CA125 as a diagnostic test in ovarian cancer: prospective validation of the Risk of Ovarian Malignancy Algorithm. *British journal of cancer* **104**(5) 863–870.

van Roosmalen, Mariëlle S., Verhoef, Lia C.G., Stalmeier, Peep F.M., Hoogerbrugge, Nicoline & van Daal, Willem A.J. 2002. Decision Analysis of Prophylactic Surgery or Screening for BRCA1 Mutation Carriers: A More Prominent Role For Oophorectomy. *Journal of Clinical Oncology* **20**(8) 2092–2100.

Vila-Parrish, Ana R., Ivy, Julie S., King, Russell E. & Abel, Steven R. 2012. Patient-based pharmaceutical inventory management: a two-stage inventory and production model for perishable products with Markovian demand. *Health Systems* **1**(1) 69–83.

Vinod, H. D. 2004. Ranking mutual funds using unconventional utility theory and stochastic dominance. *Journal of Empirical Finance* **11**(3) 353–377.

Walsh, Peter D., Biek, Roman & Real, Leslie A. 2005. Wave-like spread of Ebola Zaire. *PLoS Biology* **3**(11) e371.

Ward, Andrew. 2016. Novartis leads way for big pharma cost-cuts. URL https://next.ft.com/content/8dafdef0-c574-11e4-bd6b-00144feab7de. Retrieved on October 5$^{\text{th}}$, 2015.

White, Chelsea C. & White, Douglas J. 1989. Markov decision processes. *European Journal of Operational Research* **39**(1) 1–16.

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.

Wilson, David P., Kahn, James & Blower, Sally M. 2006. Predicting the epidemiological impact of antiretroviral allocation strategies in KwaZulu-Natal: the effect of the urban-rural divide. *Proceedings of the National Academy of Sciences of the United States of America* **103**(38) 14228–14233.

Woolhouse, Mark E. J., Rambaut, Andrew & Kellam, Paul. 2015. Lessons from Ebola: improving infectious disease surveillance to inform outbreak management. *Science translational medicine* **7**(307) 1–8.

Wooster, R., Neuhausen, S., Mangion, J., Quirk, Y., Ford, D., Collins, N., Nguyen, K., Seal, S., Tran, T., Averill, D. & et, al. 1994. Localization of a breast cancer susceptibility gene, BRCA2, to chromosome 13q12-13. *Science* **265**(5181) 2088–2090.

World Health Organization. 2003. The operational response to SARS. URL http://www.who.int/csr/sars/goarn2003_4_16/en/. Retrieved on May 5th, 2014.

World Health Organization. 2015. Ebola virus disease. URL http://www.who.int/mediacentre/factsheets/fs103/en/. Retrieved on October 5th, 2015.

World Health Organization. 2016. Ebola response funding. URL http://www.who.int/csr/disease/ebola/funding/en/. Retrieved on May 17th, 2016.

Wu, Joseph T., Riley, Steven & Leung, Gabriel M. 2007. Spatial considerations for the allocation of pre-pandemic influenza vaccination in the United States. *Proceedings of the Royal Society B: Biological Sciences* **274**(1627) 2811–2817.

Yu, Zhenxin, Yan, Hong & Cheng, T. Edwin. 2001. Benefits of information sharing with supply chain partnerships. *Industrial management & Data systems* **101**(3) 114–121.

Zaric, Gregory S. & Brandeau, Margaret L. 2001. Resource allocation for epidemic control over short time horizons. *Mathematical Biosciences* **171**(1) 33–58.

Zaric, Gregory S. & Brandeau, Margaret L. 2002. Dynamic resource allocation for epidemic control in multiple populations. *Mathematical Medicine and Biology* **19**(4) 235–255.

Declaration

Last Name: Nohdurft **First Name:** Eike

Affirmation – Statutory Declaration
According to § 10 part 1 no. 6 of the Doctoral Program Regulations
(dated 5th March 2008 as amended on the 8th March 2012)

I hereby declare that the

Dissertation

submitted to the

Wissenschaftliche Hochschule für Unternehmensführung (WHU)
- Otto-Beisheim-Hochschule -

was produced independently and without the aid of sources other than those which have been indicated. All ideas and thoughts coming both directly and indirectly from outside sources have been noted as such.

This work has previously not been presented in any similar form to any other board of examiners.

Sentences or text phrases, taken out of other sources either literally or as regards contents, have been marked accordingly. Without notion of its origin, including sources which are available via internet, those phrases or sentences are to be considered as plagiarisms. It is the WHU's right to check submitted dissertations with the aid of software that is able to identify plagiarisms in order to make sure that those dissertations have been rightfully composed. I agree to that kind of checking, and I will upload an electronic version of my dissertation on the according website to enable the automatic identification of plagiarisms.

The following persons helped me gratuitous / non-gratuitous in the indicated way in selecting and evaluating the used materials:

Last name	First name	Kind of support	Gratuitous / non-gratuitous
Rost	Jürgen	Provision of Demand data for Nohdurft & Spinler (2016)	Gratuitous
Long	Elisa	Co-author of Nohdurft et al. (2016a,b)	Gratuitous

Further persons have not been involved in the preparation of the presented dissertation as regards contents or in substance. In particular, I have not drawn on the non-gratuitous help of placement or advisory services (doctoral counsels / PhD advisors or other persons). Nobody has received direct or indirect monetary benefits for services that are in connection with the contents of the presented dissertation.

The dissertation does not contain texts or (parts of) chapters that are subject of current or completed dissertation projects.

Hamburg, February 26th, 2017

Eike Nohdurft