Three Practical Problems in Healthcare Analytics

by

Vusal Babashov

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

Telfer School of Management

(Management)

THE UNIVERSITY OF OTTAWA

(Ottawa)

February 2021

Abstract

This thesis investigates three critical problems faced by healthcare service providers and proposes analytical solutions. In the first manuscript, we aim to set wait time targets in a multi-priority patient setting using simulation, statistical regression and supervised machine learning. Using illustrative examples with two and three patient classes, and a clinical study with four patient priority classes, we demonstrate potential managerial and societal savings based on the proposed approach. Numerical analyses show that typically wait time targets are quite low, and thus the additional wait imposed by long wait time targets cannot be justified.

In the second manuscript, we investigate scheduling policies to book patients for follow-up appointments with their service providers dynamically. We develop a Markov decision process model to efficiently allocate available capacity to consults and follow-up visits in a dynamic fashion. We solve this model using the linear programming approach to Approximate Dynamic Programming (ADP) and discuss the characteristics of the approximate optimal booking (AOP) policy for multi-class patients with repeat visits. Finally, we compare the AOP policy's performance to that of existing policies through simulation and show the superior performance of the AOP policy over a First Available Slot Policy with booking limits (i.e., a Myopic policy).

In the third manuscript, we propose a framework for the drug formulary decision with the help of Multi-Criteria Decision Analysis. We use a recent extension of the UTilities Additives DIScriminantes approach, UTADIS^{GMS} and demonstrate the method using the oncology drugs reviewed through pan-Canadian Oncology Drug Review (pCODR) in Canada between 2011 and 2017. Finally, we show the method's prescriptive and predictive ability using an open-source decision support tool.

Each paper makes a worthwhile contribution to the healthcare operations literature through the modelling and case studies using real-life data. Also, each essay conclusion has the potential to inform healthcare policy decision making. Approaches presented in

the thesis can help managers allocate resources efficiently with proven, analytics-based methods.

Table of Contents

List of Tables

List of Figures

Acknowledgements

This thesis would not have been what it is without the encouragement and support from many individuals. I am grateful to every person who offered help to me and contributed in some way to the process leading to my dissertation.

It is difficult to describe in words my deep gratitude to my supervisor, Jonathan Patrick and my co-supervisor, Antoine Sauré. With their patience, inspiration, and willingness to explain things clearly and simply to me, they have been great mentors throughout my Ph.D. studies. I would like to thank both for our research meetings, advice and insightful direction.

I have also been fortunate enough to have a great thesis committee. Besides Jonathan Patrick and Antoine Sauré, I had the privilege to have Onur Öztürk, Gilles Reinhardt and Jonathan Li as my thesis committee members. I thank them for being supportive of my research and for their encouragement and constructive critiques that have greatly contributed to my thesis.

I would like to thank a few individuals who contributed to this research. I will start with Ahmet Kandakoglu, a post-doctoral research associate at Telfer, who lended to me his coding expertise at critical times. I also thank Gilles Reinhardt and Sarah Ben Amor, co-authors of Chapter 4, for introducing me to the topic of Multi-Criteria Decision Analysis (MCDA). Finally, I want to thank Davood Astaraky, a methodologist at The Ottawa Hospital, for facilitating the clinical data used for the case studies presented in Chapter 2 and Chapter 3. Furthermore, thanks to the University of Ottawa, Telfer School of Management and NSERC for their financial support during my studies.

Last but not least, I want to thank my mom, dad, and wife. I am forever indebted for their understanding, endless patience and encouragement when it was most required.

To my family

Co-Authorship Statement

Chapters 2 and 3 correspond to manuscripts co-authored with Professors Jonathan Patrick, Antoine Sauré and Onur Öztürk. The identification and design of the research program for these papers were carried out jointly with Professors Jonathan Patrick and Antoine Sauré. The research, analysis and manuscript preparation were performed by the candidate under close supervision from Jonathan Patrick and Antoine Sauré, and with comments from Onur Öztürk.

Chapter 4 is a published article co-authored with Professors Gilles Reinhardt and Sarah Ben Amor. The identification, design of the research program, manuscript preparation and revisions were carried out jointly. The model implementation and analysis were performed by the candidate.

1 Introduction

We begin by motivating three managerial and practical healthcare problems such as wait time targets, follow-up patient scheduling and drug formulary listing decisions in Section 1.1. Then, in Section 1.2, we give an overview of the thesis. Finally, in Section 1.3, we provide a detailed outline of the thesis.

1.1 Motivation

The healthcare industry continues to grow while being one of the largest service industries in developed countries around the world. In 2018, the healthcare spending constituted more than 10.7% of the GDP in Canada and 16.9% in United States, the highest in the developed world. This equates to more than $\$6,448$ per capita in healthcare spending in Canada and $\$13,722$ per capita in the US [2]. Globally, the healthcare industry faces increasing pressure to reduce costs, improve care coordination and patient outcomes together with improving both clinical and operational efficiency by providing more with less. Not only costs but also demand continue to grow both in Canada and around the world [6]. Possible remedies to deal with these challenges include increasing available capacity, managing demand and improving efficiency. In most cases, the first two alternatives are not possible leaving the third option as the only viable option. However, improving efficiency requires an extensive analysis of service operations which can be supported by a range of analytical tools and methodologies that help make informed scheduling and capacity allocation decisions.

Uncertainty can make the scheduling and capacity allocation decisions complex and challenging. Typically, the main source of uncertainty for this type of problems is the demand for care. In such an environment, one needs to find a balance between patient waiting times and the ability to efficiently use clinical resources. If clinics book patients (primarily non-urgent or lower priority patients) too soon, there might be insufficient capacity for later-arriving higher priority patients. On the other hand, if clinics book patients too far into

the future, patients might wait longer than needed and there is potential for idle capacity. Therefore, it is critical to have good and effective schedules to ensure that patients receive timely access to medical services in a cost-efficient manner. This is especially important now that healthcare costs and demand for medical services are on the rise and, consequently, health systems in many jurisdictions around the world are under constant stress to improve health care efficiency while reducing costs.

Wait times are among the most significant problems in healthcare. They are a consequence not only of an imbalance between capacity and demand but also a result of inefficient patient scheduling. Demand for health care services has increased dramatically in the past decades as a result of an aging population, longer life expectancy, pervasiveness of chronic and infectious diseases, limited financial resources to augment or replace existing capacity, and the development of new medical screening and treatment technologies. Capacity, on the other hand, has failed to keep up with demand and limited attention has been paid to identifying efficient resource allocation practices. Reducing wait times under these circumstances has become one of the most challenging problems facing healthcare providers. This is mainly due to the complexity introduced by a wide range of services, different urgency categories of patients, and multiple types of resource requirements. Clearly, the efficient allocation of medical resources to patient needs requires precise coordination and management of available resources, a task often so complex that it is beyond what a person can do manually.

The long term workload of a health provider (i.e., panel size) is impacted by new admissions and existing patients who are followed during or after treatment since patient follow-ups are performed by a same physician or care provider. If a care provider's workload is large, their availability for new patient consults is limited thus increasing the wait time for new patient appointment. New patient consults with care provider are usually scheduled on a first-come-first-served basis, without considering provider's availability. Effective patient scheduling policies should take into account these factors and consider different follow-up practices to better estimate the immediate and future impact of appointment scheduling decisions and thus provide timely access to care. Even though patient scheduling problems have been studied extensively in the literature, the dynamic allocation of medical capacity in advance of the service date in the presence of multiple follow-up visits has received limited attention.

A wait time target is defined as the maximum medically recommended wait time. For a patient to receive a care, wait time targets are usually set by government, health care organizations and medical associations upon consensus of physicians and healthcare administrators with the aim to prioritize, control and smoothen the flow of patients. Faced with insufficient resources, a common current clinical practice is to book less urgent patients further into the future by extending the wait time window. From an analytics stand point, two key elements that are critical in setting wait time targets are the resulting capacity utilization and the demand levels both of which are not considered through the consensus approach. There is a clear trade-off in how wait time targets are set. Extending targets will smooth out the demand for care over a period of time and thus reduce the need for overtime but at the same time, it will increase patient wait times and associated costs. Thus, a challenging task for a healthcare facility is to book randomly arriving patients into available capacity while meeting priority-specific wait time targets in a cost-efficient manner.

Pharmaceutical drug expenditures account for a large share of the healthcare expenditures in most countries. In 2016, drug spending accounted for $\$1,043$ per capita in Canada and $\$1,470$ per capita in United States [1]. Payers for drugs can be the public sector (federal or provincial governments) or the private sector (private health insurance companies, and individuals or households who pay out of pocket to cover drug expenses). For example, in Canada prescription drug coverage is available to the labour force from private insurance plans and to seniors, first nations and armed services members through public or government insurance plans [9]. The premium and deductible levels vary considerably between plans but, overall, provincial drug plans account for 28% to 41% of the total drug expenditure whereas private plans account for 35% of the total drug expenditure [7, 3]. Whether it is a public or private plan, the process starts with a physician writing prescriptions to patients who then purchase the prescribed drugs from pharmacies at a cost to be reimbursed fully or partially by a payer such as a public or private health insurance plan.

A medical drug formulary is a list of drugs that payers approve to reimburse. Provincial governments maintain their own public plans. In addition, cancer agencies and large hospitals in each province and city have their own drug formularies. Drugs that are in the formulary are said to be "listed". Pharmaceutical drug companies want their products listed because physicians and/or patients prefer drugs that will be reimbursed, otherwise, physicians will substitute unlisted drug with a similar and competing listed drug that will

be reimbursed by the payer.

Payers, both public and private, require manufacturing companies to submit a thorough clinical and economic analysis report about their products as a part of their formulary application[11]. Formulary listing decisions in Canada are guided by the Canadian Agency for Drugs and Technology in Health (CADTH)'s Common Drug Review (CDR) process. Subsequent to the review, the Canadian Drug Expert Committee (CDEC) provides one of the following recommendations to public drug plans: "List", "List with clinical criteria and/or conditions", "Do not list at the submitted price" or "Do not list" [5]. The CDEC considers the following established criteria in making a listing recommendation: clinical studies which detail safety, efficacy and effectiveness of the drug compared to alternatives, therapeutic advantages/disadvantages and cost-effectiveness relative to chosen comparators [4]. However, CDR recommendations are non-binding, and it is ultimately up to provincial health authorities to determine if a new drug will be covered in their plans. This results in coverage variation across plans and across provinces [8]. For private payers, there is no centralized review process, so drug manufacturers need to apply to be listed for each payer individually. Other countries such as the United States and the United Kingdom have a similar process.

Formulary design is a complex decision making problem because it involves multiple factors such as heterogeneity in patient characteristics, unit price, economic evaluation (i.e., incremental cost-effectiveness ratio or ICER), budget impact analysis, evidence on clinical benefits, safety, cost-effectiveness variation across multiple conditions, quality of clinical evidence, availability of alternative treatment options, and burden of illness, all of which factor into a final listing decision. Often there are a number of alternatives available. However, the information available about each alternative is imperfect. Since there exists conflicting trade-offs between the decision alternatives, decision makers need to prioritize what matters most to allocate the limited budgets. The cognitive burden of the process has led to the use of certain heuristics such as an ICER threshold of $\$50,000$ per quality adjusted life years (QALY) in some countries. This type of approach can lead to suboptimal decisions in formulary design [12]. The decision process becomes even more difficult if multiple individuals such as committee is involved leading to conflicting priorities between the individuals. In addition, several external factors such as public pressure by physician or patient groups and turnover in review committee members may further complicate the

process resulting in subjective and inconsistent decisions over time.

1.2 Overview of the Thesis

1.2.1 Chapter 2: Setting Patient Wait Time Targets in a Multi-Priority Patient Setting

In Chapter 2, we study wait time targets in a multi-priority patient setting. As discussed briefly above, wait time targets are typically determined through collective agreement and are not based on analytical reasoning. The patient scheduling practice, available capacity and demand for care all play an important role in achieving wait time targets but all are ignored by the consensus approach frequently used in practice. We propose a three-step approach to set wait time targets: First, a good and efficient scheduling policy must be chosen. We discuss existing booking policies and select the "The Day with the Minimum number of Bookings" (DMB) policy [10] to use in our analysis. Second, the chosen scheduling policy is simulated to estimate average wait times and overtime utilization. This step is setting-specific given that it takes a specific capacity level, demand characterizations, and scheduling policy into account. Third, we determine the implicit unit waiting cost for a given choice of wait time targets using inverse optimization. We provide several numerical examples, including a cases study based data from a rheumatology clinic. Numerical results show that typically patients wait longer than needed for no meaningful benefit to hospitals or clinics in terms of resource management. In addition, we show that under certain circumstances, the wait time targets can be reduced, with minimal additional resource requirements, enabling clinics to judiciously manage their own resources while providing patients with more timely access to medical services.

1.2.2 Chapter 3: Dynamic Advance Patient Scheduling with Follow-up Appointments

In this chapter, we study the advanced scheduling of multiple recurring visits for a single care provider. Allocating capacity to randomly arriving single appointment patient requests is already a challenging problem to solve. Extending it to the multiple visits setting is even more complex. We develop a generic Markov decision process (MDP) model in which the number of additional visits can be known in advance or occur according to a given

probability mass function. Based on the AOP for small settings, we derive heuristic policy and compare its performance against that of Myopic and DMB policies via a simulation in Java and evaluate the performance of the policy for real-world case study based on data from an endocrinology outpatient clinic. Numerical results show significant cost savings and improved service levels when the proposed heuristic policy is used in this setting instead of the Myopic policy. The difference in performance between the heuristic policy and the DMB policy is not statistically significant.

1.2.3 Chapter 4: Framework For Drug Formulary Decision Using Multi Criteria Decision Analysis

In this chapter, we study drug formulary design decisions. As mentioned above, formulary design is a complex decision making problem that involves both quantitative and qualitative criteria. We highlight some of the shortcomings of the existing drug listing decision and present the robust ordinal regression-based Multi-Criteria Decision Making (MCDA) approach known as UTADIS^{GMS} and the associated open-source decision support tool. To demonstrate the potential benefits of this method, we provide results based on published pan-Canadian Oncology Drug Review (pCODR) data.

1.3 Outline of the Thesis

The rest of this thesis is organized in a series of chapters. At the beginning of each chapter, we motivate the problem being addressed and examine the related literature. We then provide our analysis and results. We conclude each chapter with a summary of our main findings. In addition to the chapters discussed in Section 1.2, Chapter 5 summarizes the thesis contributions and provides a brief discussion of future research directions.

1.4 Bibliography

[1] Organisation for Economic Co-operation and Development. OECD Health Statistics 2018. URL `https://data.oecd.org/healthres/health-spending.htm`.

[2] Organisation for Economic Co-operation and Development. OECD Health Statistics 2019. URL `https://data.oecd.org/healthres/health-spending.htm`.

[3] Steve Barnes and Laura Anderson. *Low earnings, unfilled prescriptions: employer-provided health benefit coverage in Canada*. Wellesley Institute, 2015.

[4] CADTH. Canadian Expert Drug Advisory Committee Terms of Reference. URL `https://www.cadth.ca/media/cdr/committees/cedac/cdr_cedac_terms_reference_e.pdf`. 2008.

[5] CADTH. Common Drug Review. Procedure for the Common Drug Review. URL `https://www.cadth.ca/media/cdr/process/Procedure_for_CADTH_CDR.pdf`. 2014.

[6] Michael Carter. Diagnosis: mismanagement of resources. *OR MS TODAY*, 29(2): 26–33, 2002.

[7] Jamie R Daw and Steven G Morgan. Stitching the gaps in the Canadian public drug coverage patchwork? a review of provincial pharmacare policy changes from 2000 to 2010. *Health Policy*, 104(1):19–26, 2012.

[8] Meghan McMahon, Steve Morgan, and Craig Mitton. The common drug review: a nice start for Canada? *Health Policy*, 77(3):339–351, 2006.

[9] Steven G Morgan and Katherine Boothe. Universal prescription drug coverage in Canada: Long-promised yet undelivered. In *Healthcare management forum*, volume 29, pages 247–254. SAGE Publications Sage CA: Los Angeles, CA, 2016.

[10] Antoine Sauré, Jonathan Patrick, and Martin L Puterman. Simulation-based approximate policy iteration with generalized logistic functions. *INFORMS Journal on Computing*, 27(3):579–595, 2015.

[11] L Tilson and M Barry. European pharmaceutical pricing and reimbursement strategies. *National Centre for Pharmacoeconomics in Ireland: Dublin*, 2005.

[12] Van-Anh Truong. Optimal selection of medical formularies. *Journal of Revenue and Pricing Management*, 13(2):113–132, 2014.

2 Setting Wait Time Targets in A Multi-Priority Patient Setting: An Inverse Optimization Approach

We consider the problem of determining wait time targets in a multi-priority patient setting where each class has a pre-defined maximum recommended wait to access the service. Using simulation, statistics, machine learning and inverse optimization, we show that long wait time targets are not justified when considering the cost of wait and care provider overtime.

2.1 Introduction and Motivation

Timely access to healthcare continues to be a growing concern in jurisdictions with both publicly- and privately-funded health systems [29, 11]. It is known that long patient wait times are associated with patient distress and adverse health outcomes [12, 25]. Health policy makers have implemented different intervention strategies to mitigate the consequences of long waits for care [23, 3, 26]. However, the impact of regulatory policy interventions has been inconclusive [5]. In Canada, patient wait times have increased over the last decade. For example, the average wait time from a general practitioner visit to a specialist visit (i.e., referral) and from a specialist visit to the start of treatment across 12 medical specialties increased more than twofold across Canada from 1993 to 2017 [6].

A wait time target is typically defined as the maximum medically recommended wait patients can endure. Most countries have implemented some form of wait time target for non-emergency medical services. Here we use the example of Canada to illustrate. In 2005, the Canadian provinces and territories agreed to common wait time targets and federal benchmarks across the country in an effort to reduce patient wait times. This pan-Canadian initiative, that was claimed to be based on "research and clinical evidence", applies to the provision of non-emergency medical treatments and screening services including hip and

knee replacement, hip fracture repair, cataract surgery, cardiac bypass surgery, radiation therapy, CT scan, MRI scan, and cancer-related surgeries. The federal benchmark for each of these services provides a single wait time target. For example, the wait time target for radiation therapy to treat cancer is four weeks for patients having ready-to-treat status, for hip fracture fixation is 48 hours, and for cataract surgery is 16 weeks for high risk patients. However, individual provinces are given the flexibility to define several urgency or priority levels, with separate wait time targets, all respecting the federal benchmark. For example, the province of Ontario has set four levels of urgency to see a surgical specialist and undergo surgery. Table 2.1 summarizes the provincial government-mandated patient priorities and associated wait time targets in Ontario for a number of non-emergency medical/clinical services. Emergency patients must be seen immediately so they are not included in the benchmark indicators that provinces report to the Canadian Institute for Health Information (CIHI). Wait time targets for non-emergency medical conditions range from 24 hours to 182 days. In the case of breast cancer, patients with priorities 1, 2, 3 and 4 seeking their first surgical appointment should be seen by a doctor within 1, 10, 21, and 35 days, respectively. In addition, Cancer Care Ontario has established wait time targets for radiation therapy based on three priority categories. Patients with priorities 1, 2 and 3 and ready-to-treat status should start radiation therapy within 1, 7 and 14 calendar days, respectively (Source: Cancer Care Ontario). In addition to government-driven regulatory initiatives, the Wait Time Alliance (WTA) comprising several medical associations has determined medically acceptable wait times for 1000 treatments in 16 areas of medical practice.

What is common to all of these wait time target initiatives is the consensus approach to their determination. Set by a panel of surgeons, specialists and health care administrators, these priority-specific wait time targets are based on clinical urgency and reasonable waits for medical procedures but they do not consider the efficient use of clinical resources and the patient volume associated with each patient class. In practice, two critical elements that are not considered in setting the targets are the system capacity and the variability in demand. Ignoring these factors results in sub-optimal operational practices. Typically, the challenging task of booking randomly arriving patients into available capacity while meeting priority-specific wait time targets in a cost-efficient manner is left to the discretion of the booking clerk with little to no guidance.

From a societal perspective, the cost of waiting and the cost of overtime are two impor-

tant components of the total cost to the health system. The cost of time lost due to waiting can be seen as the monetary value of lost productive time. Many patients can experience discomfort, pain and even health deterioration resulting in lost productivity and wages. A recent Statistics Canada study estimated that the average cost of waiting for 1,082,541 patients who waited 11 weeks on average for treatment, and reported that the wait adversely affected their lives, was $1,924 per patient [7]. The above attempt notwithstanding, the cost of waiting is difficult to quantify. In this chapter, we take an implicit rather than explicit approach to estimate the cost of waiting. We recognize that the setting of wait time targets is a balancing act that seeks a compromise between short wait times and reduced overtime utilization. Extending wait time targets allows health care organizations to smooth out demand over a period of time and thus mitigate against the need for overtime. However, longer targets lead to increased patient wait times and the associated cost. Thus, we are attempting to balance a subjective cost (waiting) against a more objective one (overtime).

Through this chapter we aim to provide a *method* to determine the implicitly imposed waiting time cost associated with choosing a set of wait time targets given a known overtime cost. This can then be used to determine the most efficient choice of targets - ones that achieve the maximum reduction in overtime for the minimum increase in patient waiting time. To that end, we first derive an analytical form for the total cost function that characterizes the trade-off between average patient wait times and the use of overtime given a possible set of wait time targets. We do this through a combination of classical and deep neural network based regression. We then use an inverse optimization approach to determine the implicit cost of waiting for any set of wait time targets. We show that the total cost function does not continue to decrease as the target for each patient class is increased beyond a given threshold meaning that there is no benefit in extending the targets further. In addition, we also show that these thresholds primarily depend on the wait time target associated with the highest priority class.

The rest of the chapter proceeds as follows. In Section 2.2, we provide a review of the relevant literature on the definition and determination of wait time targets. In Section 2.3, we describe the proposed methodological approach in detail. We then discuss a number of numerical experiments in Section 2.4 and conclude with remarks and future directions in Section 2.5.

Table 2.1: Government-Mandated Wait Time Targets in Ontario

Surgeries and Procedures	**Priority 1**	**Priority 2**	**Priority 3**	**Priority 4**
Pediatric Surgeries				
Time to First Surgical Appointment	-	30 days	90 days	182 days
Cancer Surgeries				
Time to First Surgical Appointment	24 hours	10 days	21 days	35 days
Time to Cancer Surgery	24 hours	14 days	28 days	84 days
Cardiac Surgeries and Procedures				
Time to Cardiac Bypass	14 days	14 days	42 days	90 days
Time to Angioplasty	7 days	7 days	14 days	28 days
Time to Angiography	7 days	7 days	28 days	84 days
Eye Surgeries				
Time to Specialist Appointment	7 days	30 days	90 days	182 days
Time to Cataract, Cornea and other	24 hours	42 days	84 days	182 days
Time to Corneal Transplant	24 hours	14 days	42 days	112 days
Time to Vitrectomy	24 hours	7 days	42 days	84 days
Time to Combination of Cataract and Other Procedures	24 hours	7-42 days	42-84 days	84 -182 days
Orthopaedic				
Time to First Orthopaedic Appointment	-	30 days	90 days	182 days
Time to Orthopaedic Surgery	7 days	7-42 days	56-84 days	182 days
Other Surgeries				
Time to First Surgical or Specialist Appointment	-	30 days	90 days	182 days
Time to Other Surgery or Procedure	-	7-56 days	56-112 days	182 days
Diagnostic Imaging				
Time to MRI or CT scan	24 hours	48 hours	10 days	28 days

Source: `http://www.hqontario.ca` [Accessed: May 1, 2019]

2.2 Related Literature

In the patient scheduling literature, wait time targets have been used as an input parameter in mathematical programming, discrete-event simulation and queuing theory models developed to reduce wait times, increase patient throughput and maximize resource utilization. For example, Patrick et al. (2008) developed a Markov Decision Process (MDP) model to dynamically schedule patients with different priorities and associated wait time targets to a diagnostic facility [28]. Similarly, Sauré et al. (2012) built and approximately solved an MDP model to dynamically schedule multiple types of radiation therapy treatments into treatment units within the mandated wait time targets [31]. More recently, Astaraky et al. (2015) developed an MDP model to schedule multi-class surgery patients with class-specific wait time targets with the goal of optimizing the use of upstream and downstream hospital resources simultaneously [4]. Huang et al. (2015) proposed a queuing model for multi-class emergency department patients to be seen by a physician within patient-specific wait time targets [17]. The authors also identified policies that asymptotically minimize the congestion costs for this setting. Li et al. (2015) developed a mixed integer programming model to optimize LINAC (linear accelerator) capacity allocation and case-mix decisions while meeting wait time target requirements for multiple types of patients [21]. Ma et al. (2016) built simulation- and optimization-based models to enable clinic managers to make informed capacity planning decisions with respect to multi-class cancer patients seeking consult appointments with oncologists [24]. Geng and Xie (2016) considered the advance scheduling of two types of patients with specific wait time targets for a diagnostic service using a finite-horizon MDP approach and provided a characterization of an optimal control policy by proving the monotonicity and concavity properties of the reward function [14]. Kazemian et al. (2017) considered five patient priority levels and a maximum wait time target for elective surgery in their discrete-event simulation model [19]. In a stochastic programming model meant to maximize both revenue and equity, Zhou et al. (2018) incorporated access time targets for different types of patients as a constraint in the allocation of hospital ward capacity [34].

There are only a few published studies in the medical research literature concerned with the *determination* of wait time targets. Paterson et al. (2006) defined the maximal wait times for specialist consultation and procedures for digestive diseases [27]. Their approach

was based on seeking a consensus on predetermined statements on medically acceptable wait times among the 25 members of a consensus group comprising community and academic gastroenterologists from across Canada. As a result, the Canadian Association of Gastroenterology defined four acuity patient categories and associated targets - within 24 hours, two weeks, two months, and six months.

Schaafsma (2006) pointed out the need for a better algorithm and reproducible methodology to determine wait time targets as opposed to the survey-and-consensus method [32]. He proposed a method that is based on the cost minimization problem that arises from the trade-off between excess capacity and waiting in a setting where random patient arrivals occur. The method addressed two elements that the consensus method did not consider: a variable number of patient arrivals and capacity allocation practices. Harrison and Appleby (2009) criticized the national 18-week wait time target set by NHS England and argue in favour of variable wait time targets across medical conditions, regions, and hospitals [15]. They argued that patient prioritization would generate more benefits without any additional costs, or a reduction in the maximum mandated wait, while also reducing patient wait times further. In a related paper, Harrison and Appleby (2010) conducted a cost-benefit analysis of different measures to reduce waits and discussed the factors that should be considered in estimating the cost of waiting and shortening wait lists [16]. They also highlighted that the need for more capacity was overestimated by the government when setting the 18-week target for a range of medical services, suggesting that waiting times can be reduced through better process flow, diversion of care from hospitals, pooling of consultants within specialties, and some rationing.

Ahn et al. (2011) empirically derived maximal acceptable wait times for adolescent scoliosis treatment [1]. Using retrospective clinical data about 216 patients, they built a logistic regression model to investigate the relationship between surgical wait times and predetermined adverse events while controlling for several confounding variables. They showed that longer wait times are associated with higher odds of adverse events. They also looked at true and false positives for adverse events at cut-off wait times ranging from 1 to 365 days. Through a Receiver Operating Characteristic curve (ROC), they showed that an access target of three months compared to a consensus target of six months has the potential to reduce the risk of adverse complications and that there was no potential gain, and hence costs savings, in setting longer targets to reduce false positives.

The most related paper in the patient scheduling domain is the one by Liu where he studied the impact of the size of the appointment scheduling window (defined by the wait time target) on operational efficiency under patient no-shows. Using a single-server queuing model, he proposed an approach to determine the capacity of the queue, which serves as a proxy for the length of the appointment scheduling window, and provided an analytical characterization for the corresponding optimal values. Our study significantly differs from that of Liu's. We consider multiple classes of patients defined in terms of clinical urgency whereas he considered only two classes of patients defined based on no-show probabilities. Furthermore, it is not clear if his conclusions can be generalized to three or more patient classes.

2.3 Methodology

There is a clear relationship between the choice of wait time targets and the use of overtime. On the one hand, long wait time targets allow patient booking clerks to book arriving demand further into the future smoothing out the variation over time. This reduces overtime usage but increases patient wait times. On the other hand, short wait time targets limit the agents' ability to book incoming patients within the targets without resorting to overtime.

In this section, we develop a parametric mathematical programming model that considers the trade-off between average patient wait times and daily overtime utilization. Importantly, we assume that the scheduling policy is known in advance. Thus, given cost parameters and the scheduling policy, the model finds a set of wait time targets (restricted by a medically acceptable upper bound) such that the total cost associated with the resulting average wait times and daily overtime utilization is minimized. In mathematical terms, the model can be formulated as follows:

$$\min_{\vec{T}\in\mathbb{Z}_+^I} \left\{ c_1 f_1^{WT}(\vec{T}) + c_2 f_2^{WT}(\vec{T}) + \ldots + c_I f_I^{WT}(\vec{T}) + h f^{OT}(\vec{T}) : T_{i-1} \leq T_i \leq T_i^{Max} \ \forall i \in I \right\}, \tag{2.1}$$

where c_i is the unit (i.e., daily or weekly) wait time cost for priority class i patients, h is the overtime cost per patient, and I is the number of patient classes. The functions $f_i^{WT}(\vec{T})$

and $f^{OT}(\vec{T})$ represent the average wait time for priority class i patients and the average daily (or weekly) overtime utilization, respectively. The vector $\vec{T} = (T_1, T_2, \ldots)$ is a decision variable that represents the wait time targets and $\vec{T}^{Max} = \left(T_1^{Max}, T_2^{Max}, \ldots\right)$ is the vector of maximum clinically recommended wait time targets. We assume that lower values of i represent higher priority patients and thus the targets should be non-decreasing in i.

The use of *average* wait times as the metric perhaps needs some justification as, quite frequently, health organizations prefer to focus on the 90th percentile. We focus on the average since we are attempting to balance the "cost" of wait against the cost of overtime. Thus, we are interested in the wait time of the average patient rather than the outlier represented by the 90th percentile. Moreover, we focus on a scheduling policy that does not book patients late meaning that the wait time targets act like hard upper limits on the wait time of patients thus providing a bound on the 90th percentile. Nonetheless, in the simulation, we present the 90th percentile wait times to demonstrate the performance.

There are several challenges that prevent us from directly solving Equation (2.1). First, average wait times and overtime utilization depend on the demand, the system capacity and the scheduling policy employed. Second, even if those are known, there is no closed form solution for either the average wait times or the overtime functions. Finally, even if there were a function that accurately captured average wait time, there are no objective values for the unit wait time costs. In the following subsections, we address each of these challenges individually. First, we briefly describe some advance patient scheduling policies available in the literature that have previously been shown to perform well. From these policies, we choose a single booking policy to use for the rest of the paper. The choice of this policy is justified below but the method itself could be utilized based on any other policy. We emphasize that what we are providing in this research is a *methodology* that determines wait time targets that are *setting specific*. An organization using a different scheduling policy would simply need to adjust the simulation component of the methodology accordingly. Second, we simulate the proposed scheduling policy and, approximate based on the results, the average wait time and overtime functions using regression or neural networks. Finally, we develop an inverse optimization model to determine the implicit wait time costs associated with the choice of a given set of targets rather than arbitrarily imposing such costs. This analysis allows the manager to better understand the trade-off between overtime and wait time that is implicitly being navigated when targets are set.

Further it allows the manager to set targets that achieve the near-maximum reduction in overtime utilization without excessively increasing wait times. We illustrate these steps using simple examples and a real clinical case study.

2.3.1 Choosing a Scheduling Policy

Advance patient scheduling plays a critical role in the optimal allocation of capacity among different types of incoming patients. There are several advance scheduling policies discussed in the literature. Using a first available slot policy (i.e., a myopic policy based on the cost function), patients are booked as soon as possible in order of urgency. This policy fails to consider the impact of today's decision on the future performance of the scheduling system. Patrick et al. (2008) derived a set of booking rules, known as the PPQ (Patrick, Puterman and Quayrenne) policy to efficiently book incoming demand into available appointment slots [28]. According to this policy, the highest priority patients are booked as soon as possible within the wait time target and lower priority patients are booked on day one, then on the last day within the target window and finally working backwards to day two. If there is insufficient capacity within the target window, overtime is used rather than booking late. More recently, Sauré et al. (2015) compared the performance of several advance booking policies including a novel SS policy (an approach using a S-Shaped Value Function Approximation), derived from a simulation-based approximate dynamic programming model, the PPQ policy, a Myopic policy, and a heuristic policy called "Day with the Minimum Number of Bookings" or DMB policy [30]. The DMB policy outperformed the Myopic policy in terms of discounted cost, average wait times for high and medium priority patients, average time to first available appointment slot, and percentage of late bookings, and performed slightly better than the PPQ policy based on discounted cost, average wait times for medium and low priority patients, capacity utilization, and average time to first available appointment slot. However, the Myopic policy resulted in slightly higher, though not statistically significant, capacity utilization and lower diversions than the DMB policy. One can argue that the selection of the policy can also be based on the capacity utilization or percentage of diversions criterion. The expected total discounted cost can be viewed as a composite performance metric as it explicitly captures the cost associated with the diversions or overtime usage and late bookings. Thus, we use the discounted cost as the primary criterion to compare the performance of the policies in this study. Based on this

metric, a use of the Myopic policy resulted in approximately sixfold increase in long-run cost to a clinic under the study. While the SS policy performs slightly better than the both DMB and PPQ policies, unfortunately, the SS policy is setting dependent and often computationally intractable making the DMB policy more practical and easy to implement. According to the DMB policy, patients are booked in increasing priority order on the day associated with the minimum number of bookings within the target window for each respective priority class. If the capacity within the target window is insufficient, overtime is used. The DMB policy aims to balance capacity utilization between days thereby minimizing the possibility of wasted capacity. Considering both performance and practicality, we choose to simulate the DMB policy to approximate the average wait time and overtime as a function of different wait time targets though again the approach can be used with any policy. To demonstrate the impact of this change, we provide one example where the first available slot policy (with booking limits) is used instead. We illustrate the approach using several examples and then using data from a rheumatology clinic at a local hospital in Ontario.

2.3.2 Approximation of the Wait Time and Overtime Functions

2.3.2.1 Simulation of the DMB Scheduling Policy: Illustrative Examples and Clinical Case Study

We describe this stage of the proposed approach using several multi-priority patient settings including two illustrative examples with two and three patient classes and a clinical case study with four patient classes. For the illustrative examples, we consider a small clinic with an average daily demand of 10 appointment requests per day, regular-hour capacity of 10 appointment slots per day, and sufficiently high overtime capacity to serve all the demand not booked through regular-hour capacity. First, we assume that the clinic classifies incoming patients into two categories, priority 1 and 2, with maximum clinically recommended wait time targets of 7 and 14 days (i.e., $\vec{T}^{max} = [7, 14]$). The daily demand is assumed to follow Poisson distribution with means of 7 and 3 requests for each priority class, respectively. In the second illustrative example, we consider a clinic which classifies incoming demand into three patient classes with maximum clinically recommended wait time targets of 7, 14, and 21 days (i.e., $\vec{T}^{max} = [7, 14, 21]$). We assume that the demand from each patient class follows a Poisson distribution with means of 5, 3, and 2 appointment requests per day,

respectively. It is important to note that we intentionally chose decreasing demand rates for these two illustrative examples as the demand rates increase for lower priority patients in the clinical case study we describe next. This will enable us to explore the impact of both increasing and decreasing demand patterns on the selection of wait time targets.

The clinical case study is based on data from a rheumatology clinic at a large teaching hospital in Ontario. The clinic receives an average of 39 new consult patients per week from patients classified into five priority classes. The wait time targets for each class are 2, 9, 13, 26 and 52 weeks, respectively. To simplify the simulation analysis of the DMB scheduling policy, we combined the demand for consults coming from the last two classes and assumed a joint wait time target of 26 weeks (*i.e.*, $\vec{T}^{max} = [2, 9, 13, 26]$). Consequently, we consider weekly Poisson service request rates of 0.78, 7.41, 13.26 and 17.55 for patients of priority 1, 2, 3 and 4, respectively.

The simulation of the DMB policy uses a warm-up period, T_0, of 250 days and a simulation length of $T = 1250$ days. Results are averaged over $R = 30$ replications. We tested the performance of the booking policy over all possible wait time target values lower than the maximum medically recommended wait time targets and computed the following output measures: average wait time (in days or weeks) per patient for each class of patients and average volume of overtime bookings (per day or week) aggregated across all priority classes. The simulation analysis was implemented in Java installed on a Mac computer with 2.4GHz Intel Core i7 and 16GB RAM memory. The pseudo-code of the simulation procedure for the two patient class example is described below:

Figures 2.1 displays the results of the simulation of the DMB policy for two classes (i.e., base case). In this case, the mean wait times increase linearly with T_1 and T_2 as shown in Figures 2.1a and 2.1b. According to Figure 2.1b, the rate of increase of the mean wait for priority 2 patients increases as a function of T_2 and achieves its lowest value when $T_1 = 1$. This intuitively makes sense as a lower target for priority 1 patients means that priority 1 patients are served more frequently through overtime leaving more available capacity for priority 2 patients. Figures 2.1c depicts the non-linear relationship between the mean number of overtime bookings per day and the wait time target for the priority 2 class, for each value of T_1. The graph shows a sharp drop in mean overtime as targets are increased (particularly for the highest priority class) followed by an almost constant mean overtime level after the targets reach a certain point. This confirms the suspicion that there is a

Algorithm 1 Simulation Procedure, Two Patient Class Example

Require: $T_1^{Max} \geq 0$ **and** $T_2^{Max} \geq 0$, where $T_1^{Max} \leq T_2^{Max}$
  $T_0 \leftarrow$ length of simulation warm-up period
  $T \leftarrow$ length of simulation run for statistics collection
  $R \leftarrow$ number of replications
  $\lambda_1, \lambda_2 \leftarrow$ demand rates
  for $i \leftarrow 1$ to T_1^{Max} **do**
    for $j \leftarrow i+1$ to T_2^{Max} **do**
      $r \leftarrow 0$
      repeat
        for $t \leftarrow 1$ to $T_0 + T$ **do**
          Generate Poisson demand with rates λ_1 and λ_2 rates on each t
          Choose the first day with a minimum number of bookings (DMB Policy)
        end for
        return daily average wait time and total overtime bookings over days $[T_0+1, T]$
        $r \leftarrow r+1$
      until $r = R$
    end for
  end for

diminishing rate of return in terms of mean overtime reduction as the wait time targets are increased. Simulation graphs for the other settings are omitted since they provide similar conclusions.

2.3.2.2 Approximation of the Average Wait Time and Overtime Functions

Having developed a simulated data set, we use statistical regression and a deep neural network method called a feed-forward Multi-Layer Perceptron (MLP) model to obtain functional representations of the relationship between wait time targets, $\vec{T}$, and average wait times, $f_i^{WT}(\vec{T})$, and overtime, $f^{OT}(\vec{T})$. There is a trade-off between model performance and interpretability with respect to selection. On the one hand, statistical regression approaches are intuitive and can result in good predictive performance when the underlying parametric distributional assumptions are satisfied. On the other hand, deep neural network approaches including the feed forward MLP models are less intuitive and interpretable (i.e., black box) but more capable of capturing the complex non-linear relationship between input and output variables and without requiring any parametric distributional assumptions. However, there are several challenges associated with the use of the deep neural network approaches. First, a large sample of observations is required to train, and validate the model and to test its predictive performance. Second, there is a risk of model over-fitting with the inclusion

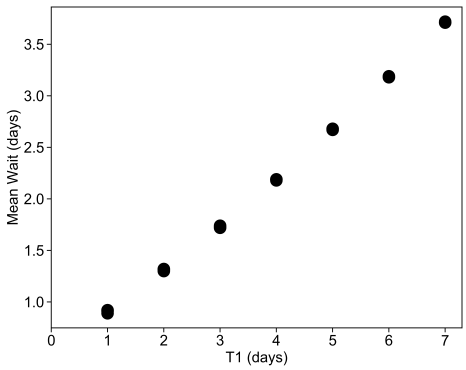

(a) Mean wait time for priority 1 patients for different values of T_1

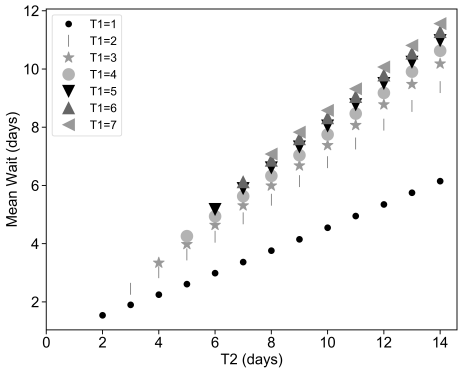

(b) Mean wait time for priority 2 patients as T_2 increases for different values of T_1

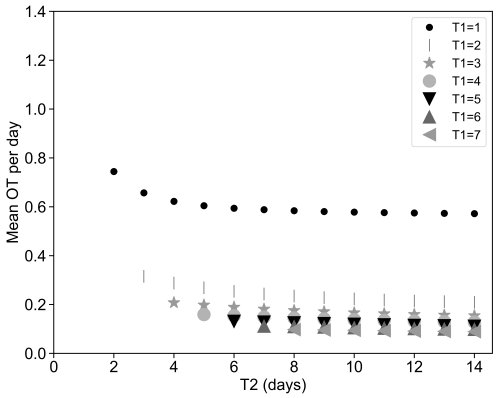

(c) Mean number of overtime (OT) bookings as T_2 increases for different values of T_1

Figure 2.1: Simulation of the DMB Policy For Two Patient Classes

of each additional neuron or hidden layer (i.e., hyper-parameters).

To obtain closed-form functions to approximate average wait times (in days or weeks) and the average number of overtime bookings (per day or week) as a function of wait time targets, we follow this step-wise procedure. First, we fit multiple linear regression, and if the goodness-of-fit measure, R^2_{adj}, is satisfactory, we stop. Otherwise, we expand the multiple linear regression model to include quadratic terms to incorporate non-linearity. Finally, we proceed to fit a feed forward MLP model if satisfactory predictive performance cannot be obtained with the quadratic and linear multiple regression models. By following this approach, we add an additional layer of complexity to the model until we get a reasonable approximation (i.e., $R^2_{adj} \geq 80\%$). In addition to quadratic terms, we also tested two-

way interaction terms between the wait time targets for different classes in a full multiple regression model. Although, some interactions terms were statistically significant (at $\alpha = 0.05$), the models did not show a meaningful increase in the R^2_{adj}. Further analysis showed that this mild interaction was only present for high values of the wait time target for the most urgent patient class. Because of the difficulties in drawing a meaningful interpretation for the cost coefficients, we omitted the two- or three- way interaction terms in the final regression models. We use R statistical software to conduct all least squares multiple linear and quadratic regression analyses. The resulting model parameters and performance results for the two, three and four patient class settings are summarized in Table 2.2. Where the statistical regression model was adequate, Table 2.2 provides the model. Where Table 2.2 indicates "Multi-Layer Perceptron Model", we resorted to a neural network approach due to the poor fit of the regression. In addition, and for comparative purposes, we present the results from two variants of the 2 patient class setting. In the first variant, patients are scheduled using the first available slot policy rather than the DMB. In the second variant, the system is assumed to be under-capacitated (imbalanced) with a capacity for nine appointments per day (but demand remaining at an average of 10 appointment requests per day). Using the regression models, the R^2_{adj} values for the predicted average wait times were over 90% for all cases except for the mean wait time for priority 1 patients in the four patient class setting. However, the R^2_{adj} for the predicted average number of overtime bookings as a function of the wait time targets was typically lower. We therefore implemented neural network models to estimate overtime.

We use Python programming language and Keras/Tensorflow machine learning framework library to fit MLP models. One of the key elements of the deep neural networks is hyper-parameter tuning. There are many techniques available to tune hyper parameters. Using an extensive manual experimentation approach, we determined the minimal number of neurons, epochs, and layers, and the batch size that provide reasonable model validation and prediction performance. We used Rectified Linear Unit (ReLU) activation functions at each input and hidden layer and a linear activation function at the output layer. (It is crucial to maintain linearity in order to be able to accomplish the next step of solving the inverse optimization model.) Results of the MLP model parameters are summarized in Table 2.3.

Using the parametric cost function form in (2.1), we can now formulate the forward

Table 2.2: Mean Wait Times and Mean Overtime Function Parameters Obtained Using Statistical Regression

Outcome	Intercept	T_1	T_2	T_3	T_4	T_1^2	T_2^2	T_3^2	T_4^2	R^2_{adj}
Two Patient Class (Base-case)										
MWT, Pr. 1	0.436	0.474	−0.007							99.1%
MWT, Pr. 2	−0.982	0.512	0.630							93.8%
MOPD	1.100	−0.285	−0.066			0.028	0.003			83.0%
Two Patient Class (First Available Slot)										
MWT, Pr. 1	1.378	0.542								99.9%
MWT, Pr. 2	1.970	0.539								99.9%
MOPD	0.235	−0.051				0.004				99.1%
Two Patient Class (Imbalanced System)										
MWT, Pr. 1		0.828	−0.014							98.7%
MWT, Pr. 2	−0.877	0.357	0.802							94.5%
MOPD				*MLP Model*						83.57%
Three Patient Class										
MWT, Pr. 1	0.574	0.381	−0.013							97.8%
MWT, Pr. 2	−0.582	0.439	0.557	−0.012						93.8%
MWT, Pr. 3	−1.578	0.372	0.263	0.647						96.2%
MOPD	0.774	−0.135	−0.058	−0.012		0.014	0.003	0.0003		79.9%
Four Patient Class (Rheumatology Clinic)										
MWT, Pr. 1				*MLP Model*						97.78%
MWT, Pr. 2	0.701		0.215	−0.003	−0.003					95.08%
MWT, Pr. 3	−0.524	0.058	0.215	0.497	−0.020					96.8%
MWT, Pr. 4	−2.230	0.092	0.117	0.212	0.848					99.3%
MOPW				*MLP Model*						93.75%

Notes: Pr: Priority, MWT: Mean Wait Time, MOPD: Mean Overtime Per Day, MOPW: Mean Overtime Per Week.

Table 2.3: Scaled Multi-Layer Perceptron (MLP) Model Parameters

Input Layer					**Output Layer**	
Weights				**Bias**	**Weights**	**Bias**
MOPD for Two Patient Class Model (Imbalanced)						
-0.475	-0.557			0.172	0.961	0.029
-0.452	0.004			0.075	1.122	
-0.466	-0.014			0.084	0.839	
-0.024	-0.036			0.000	-0.526	
-0.585	-0.036			0.134	0.945	
0.131	-0.004			0.084	-0.162	
MWT, Pr 1 for Four Patient Class Model						
0.098	0.193			1.398 2.158	0.094	0.097
-0.124	-1.095				0.115	
-0.007	0.005					
0.045	-0.032					
MOPW for Four Patient Class Model						
0.004	-0.071	-0.004		0.001	-0.111	0.037
0.039	-0.316	-0.056		0.274	2.606	
0.102	-1.302	-0.161		0.085	1.171	
0.166	-3.566	-0.136				

optimization problem in standard form denoted by "FO". The corresponding model uses the parametric functions for the average wait times and the overtime bookings determined above. The constraints ensure that the wait time targets are non-decreasing with respect to the priority class and that they respect the maximum clinically acceptable wait time for each group. The generic forward optimization model is given below

$$\mathbf{FO}(\vec{c}) \quad \min_{(\vec{T})} \quad f(\vec{c},\vec{T}) = c_1 f_1^{WT}(\vec{T}) + c_2 f_2^{WT}(\vec{T}) + c_3 f_3^{WT}(\vec{T}) + \ldots + h f^{OT}(\vec{T}) \tag{2.2}$$

subject to:

$$g_m(\vec{T}) \leq 0 \quad m = 1, \ldots, 2|I| + |I| - 1 \tag{2.3}$$

where c_i is the unit wait time cost for patient class i, $f_i^{WT}(\vec{T})$ is the average wait time function for patient class i, $f^{OT}(\vec{T})$ is the average overtime bookings, $g_m(\vec{T})$ is the mth constraint, and $|I|$ is the cardinality of the patient class set I.

Example 1: Two Patient Class Example (Base-case), Statistical Regression Assuming a fixed unit cost of $h = 100$ for each overtime booking, the explicit formulation for the two patient class setting is given below.

$$\begin{aligned}\mathbf{FO}(\vec{c}) \quad \min_{T_1,T_2} \quad & c_1\left(0.436 + 0.474T_1 - 0.007T_2\right) + c_2\left(-0.982 + 0.512T_1 + 0.630T_2\right) + \\ & +100\left(1.100 - 0.285T_1 - 0.066T_2 + 0.028T_1^2 + 0.003T_2^2\right)\end{aligned} \tag{2.4}$$

subject to:

$$T_1 \leq T_2 \tag{2.5}$$

$$T_1 \leq 7 \tag{2.6}$$

$$T_2 \leq 14 \tag{2.7}$$

$$T_i \geq 0 \quad i = 1, 2 \tag{2.8}$$

Example 2: Two Patient Class Example (Imbalanced System), MLP To illustrate the approach when neural networks are required, we provide a description of

the imbalanced variant of the two patient class setting. The ReLU activation function, $y = f(x) = Max(x, 0)$, used to derive the approximation functions for the average wait time and the overtime using the MLP model is continuous everywhere but not differentiable at $x = 0$, which makes the forward optimization model non-linear and the direct application of the Karush-Kuhn-Tucker (KKT) conditions (required in the inverse optimization framework discussed later) impossible. Therefore, we first linearize the ReLU functions and transform the piece-wise linear functions generated via the neural network approach into the equivalent linear constraints. The training performance of the deep neural network models is superior if predictors (or features) and outcome (or target) variables are normalized. We use the min-max scaler function available in the scikit-learn library in Python to scale values between zero and one. Weight and bias parameters reported in Tables 2.3 are obtained using *scaled* wait time target values, average wait times and overtime. Thus, it is necessary to reverse the transformation at the end.

In the two patient class setting with an imbalanced system, the mean overtime per day (MOPD) is approximated using a feed-forward MLP with six neurons in the input layer leading to six pairs of constraints in the linear program. Further, the MOPD function is inverse transformed before inserting it into the objective function.

$$\mathbf{FO}(\vec{c}) \quad \min_{T_1,T_2} \quad c_1\,(0.828T_1 - 0.014T_2) + c_2\,(-0.877 + 0.357T_1 + 0.802T_2) + \\ +100\,(0.78(0.961y_1 + 1.222y_2 + 0.839y_3 - 0.526y_4 + 0.945y_5 - 0.162y_6 + 0.029) + 1.01) \tag{2.9}$$

subject to:

$$-0.079(T_1 - 1) - 0.043(T_2 - 1) + 0.172 \leq y_1 \tag{2.10}$$

$$-0.075(T_1 - 1) + 0.0003(T_2 - 1) + 0.075 \leq y_2 \tag{2.11}$$

$$-0.078(T_1 - 1) - 0.001(T_2 - 1) + 0.084 \leq y_3 \tag{2.12}$$

$$-0.004(T_1 - 1) - 0.003(T_2 - 1) + 0.000 \leq y_4 \tag{2.13}$$

$$-0.098(T_1 - 1) - 0.003(T_2 - 1) + 0.134 \leq y_5 \tag{2.14}$$

$$0.022(T_1 - 1) - 0.0003(T_2 - 1) + 0.084 \leq y_6 \tag{2.15}$$

$$T_1 \leq T_2 \tag{2.16}$$

$$T_1 \leq 7 \tag{2.17}$$

$$T_2 \leq 14 \tag{2.18}$$

$$T_i \geq 0 \qquad i = 1, 2 \tag{2.19}$$

$$y_k \qquad k = 1, 2, \ldots 6 \tag{2.20}$$

Forward optimization formulations for the other settings are derived in much the same manner, and are omitted here due to size.

2.3.3 Inverse Optimization

The above sections provide an approach to determine reasonable parametric functions for the average wait time for patients of each priority class and for the number of overtime bookings for a given demand, system capacity and scheduling policy. However, the cost parameters associated with patient waiting, $c_1, ..., c_I$ remain difficult to quantify. Rather than attempting to quantify them, we use inverse optimization to determine the implicit cost of waiting suggested by a given choice of wait time targets.

Inverse optimization seeks to determine the cost parameters associated with the forward problem that render a given feasible solution optimal or approximately optimal. Inverse optimization has a wide range of applications including but not limited to finance [8], healthcare decision making [13], cancer treatment planning [10], and energy markets [9]. Ahuja and Orling discuss the general methodological framework for inverse problems under L_1 and L_∞ norms [2]. Iyengar and Kang (2005), Zang and Xu (2010), and more recently, Keshavarz et al. (2011) provided an inverse optimization framework for convex programs [18, 33, 20]. Their inverse optimization approaches are based on satisfying the KKT nec-

essary and sufficient optimality conditions with some positive residual (epsilon). In other words, the ϵ approach determines the cost parameter values that make the observed solution approximately satisfy the KKT optimality conditions [10, 8, 20].

We formulate the inverse optimization models, denoted by "IO", using the inverse optimization framework based on the necessary and sufficient KKT conditions [20]. The generic IO model for a multiple patient class setting is presented below.

$$\mathbf{IO}(\vec{T}) \quad min \ \|\epsilon\|_1 = \sum_{i,m} |\epsilon_i| + |\epsilon_m| \tag{2.21}$$

subject to

$$\nabla f(\vec{c}, \vec{T}) + \sum_m \lambda_m \nabla g_m(\vec{T}) + \epsilon_i = 0 \qquad i = 1, \ldots |I| \tag{2.22}$$

$$\lambda_m g_m(\vec{T}) + \epsilon_m = 0 \qquad m = 1, \ldots, 2|I| + |I| - 1 \tag{2.23}$$

$$g_m(\vec{T}) \leq 0 \qquad m = 1, \ldots, 2|I| + |I| - 1 \tag{2.24}$$

$$\lambda_m \geq 0 \qquad m = 1, \ldots, 2|I| + |I| - 1 \tag{2.25}$$

$$c_i \geq c_{i+1} \qquad i = 1, \ldots |I| - 1 \tag{2.26}$$

$$c_i \geq 0 \qquad i = 1, \ldots I \tag{2.27}$$

In the inverse optimization model, IO($\vec{T}$), the elements of the vector $\vec{T}$, denoted by T_i, represent observed candidate solutions. Constraints (2.22) correspond to the KKT stationarity conditions obtained by making the gradient of the Lagrangian with respect to each element of vector $\vec{T}$ equal to zero, where λ_m is a Lagrangian dual variable (or multiplier) associated with the mth constraint in the primal problem. Constraints (2.23) correspond to the complementary slackness conditions. To ensure dual feasibility of the inverse models, we include residual terms, unrestricted in sign, in the stationarity and complementary slackness conditions that are minimized in the objective. Constraints (2.24) correspond to the primal feasibility conditions and constraints (2.25) correspond to the dual feasibility conditions. In addition, in constraints (2.26) - (2.27), we assume that the elements of a feasible wait time cost vector, $\vec{c}$, are monotonic in relation to one another and that unit wait time costs are non-negative. It is important to note that although we impute the objective function based

on integer wait time target values, we do not constrain the decision variables in the above forward and inverse optimization models to be integer since non-integer (i.e, continuous) wait time target values have practical interpretation too.

Example 3: Inverse Optimization Model for the Two Patient Class (Base-case) Setting Using Statistical Regression The example below is the explicit formulation of the IO model for the case of two classes with the FO objective derived using statistical regression.

$$\mathbf{IO}(\vec{T}) \quad min \ \|\epsilon\|_1 = \sum_{m=1}^{7} |\epsilon_m| \tag{2.28}$$

subject to:

$0.474c_1 + 0.512c_2 - 28.5 + 5.6T_1 + \lambda_1 + \lambda_2 - \lambda_4 + \epsilon_1 = 0$

$-0.007c_1 + 0.630c_2 - 6.6 + 0.6T_2 - \lambda_1 + \lambda_3 - \lambda_5 + \epsilon_2 = 0$

$\lambda_1 (-T_2 + T_1) + \epsilon_3 = 0 \qquad \lambda_2 (-7 + T_1) + \epsilon_4 = 0 \qquad \lambda_3 (-14 + T_2) + \epsilon_5 = 0$

$\lambda_4 (-T_1) + \epsilon_6 = 0 \qquad \lambda_5 (-T_2) + \epsilon_7 = 0$

$\lambda_j \geq 0 \qquad \forall j$

$c_i \geq c_{i+1} \qquad i < I$

FO($\vec{c}$) constraints in Example 1

Example 4: Inverse Optimization Model for the Two Patient Class (Imbalanced System) Setting Using MLP The IO for the case of two classes with an MLP-derived FO objective function is given below. Since a total of six neurons were used in the input layer of the MLP, twelve additional constraints are required.

$$\mathbf{IO}(\vec{T}) \quad min \ \|\epsilon\|_1 = \sum_{m=1}^{14} |\epsilon_m| \tag{2.29}$$

subject to:

$$0.828c_1 + 0.357c_2 + \lambda_1 + \lambda_2 - \lambda_4 - 0.079\lambda_6 - 0.075\lambda_7 - 0.078\lambda_8 - 0.004\lambda_9 - 0.098\lambda_{10} + 0.022\lambda_{11} + \epsilon_1 = 0$$

$$-0.014c_1 + 0.802c_2 - \lambda_1 + \lambda_3 - \lambda_5 - 0.043\lambda_6 + 0.0003\lambda_7 - 0.001\lambda_8 - 0.003\lambda_9 - 0.003\lambda_{10} - 0.003\lambda_{11} + \epsilon_2 = 0$$

$$74.958 - \lambda_6 - \lambda_{12} + \epsilon_3 = 0 \qquad 87.516 - \lambda_7 - \lambda_{13} + \epsilon_4 = 0 \qquad 65.442 - \lambda_8 - \lambda_{14} + \epsilon_5 = 0$$

$$-41.028 - \lambda_9 - \lambda_{15} + \epsilon_6 = 0 \qquad 73.71 - \lambda_{10} - \lambda_{16} + \epsilon_8 = 0 \qquad -12.636 - \lambda_{11} - \lambda_{17} + \epsilon_9 = 0$$

$$\lambda_1 (-T_2 + T_1) + \epsilon_{10} = 0 \qquad \lambda_2 (-7 + T_1) + \epsilon_{11} = 0 \qquad \lambda_3 (-14 + T_2) + \epsilon_{12} = 0$$

$$\lambda_4 (-T_1) + \epsilon_{13} = 0 \qquad \lambda_5 (-T_2) + \epsilon_{14} = 0$$

$$\lambda_j \geq 0 \qquad \forall j$$

$$c_i \geq c_{i+1} \qquad i < I$$

FO($\vec{c}$) constraints in Example 2

The coefficients in the above inverse optimization problems are obtained by applying the KKT conditions to the primal or forward optimization problems in Example 1 and Example 2. The inverse optimization models for the other cases are derived in much the same manner. Due to size, the full models for the three and four patient classes are omitted.

2.3.4 Convexity

The validity of the above models depends on the convexity of the mathematical programs. In this section, we demonstrate that convexity.

Proposition 1. Primal objective functions that are affine, quadratic or a combination of both are convex in the wait time targets.

Proof. The proof is trivial. Using the convexity properties of affine and quadratic functions with positive quadratic terms, one can show that the average wait time and overtime functions are convex in the wait time targets. In addition, a non-negative weighted ($i.e., c_i \geq 0$) sum of convex functions is also convex resulting in an overall convex cost function. Alternatively, using the second order necessary and sufficient conditions of convexity, one can easily show that primal objective functions are convex. In other words, the Hessian matrix is positive semi-definite. □

Remark 2.3.1. *Due to the convexity property of norms, the inverse objective functions associated with the inverse optimization problems are also convex.*

Remark 2.3.2. *Both primal and dual problems satisfy Slater's constraint qualifications. Under Slater's rule, the KKT conditions are both necessary and sufficient for optimality.*

Corollary 2.3.3. *It is important to note that Proposition 1 and Remarks 2.3.1 and 2.3.2 ensure that a local optimal solution to the non-linear forward and inverse models will also be a global optimal solution.*

In the following section, we discuss some numerical results for the inverse convex optimization problems.

2.4 Numerical Results

The objective function of the IO problems for all models is to minimize the $\|\epsilon\|_1$ or L_1-norm of vector ϵ, also known as the Manhattan norm. Since the absolute value function is a convex non-linear function, we replaced ϵ_i by $\epsilon_i^+ - \epsilon_i^-$ (where $\epsilon_i^+ \geq 0$, $\epsilon_i^- \geq 0$) in all mathematical programs.

We implemented the inverse optimization models in Python and solved them using Gurobi API on a Mac computer with a 2.4GHz Intel Core i7 processor and 16GB RAM memory. For practical reasons, we considered all integer feasible solutions for wait time targets.

2.4.1 Two Priority Class Setting with its Two Variants

For the two patient class base-case model, Figure 2.2a demonstrates that the unit wait time cost for priority 1 patients, c_1, decreases monotonically as the value of T_1 increases and that it quickly drops to zero past $T_1 = 5$ across different values of T_2. For a fixed value of T_1, the cost coefficient c_1 attains its lowest value when $T_1 \geq 6$ and its highest when $T_1 = 1$. At its peak, the cost of an additional day of wait for priority 1 patients is almost half the cost of an overtime slot and the cost of wait for priority 2 patients is zero.

Figure 2.2c provides more insight into the unit wait time costs from a different angle. The highest value for the cost of wait for the second priority class is one tenth of the cost of an overtime slot and occurs when $T_1 = 1$ and $T_2 = 2$. As the wait time target T_2 increases, c_2 decreases monotonically hitting the cost of zero at $T_2 = 11$ for T_1 values ranging from 1 to 6 days. Finally, unit costs are zero when T_1 and T_2 are both greater than or equal to 6 days. Figures 2.2b and 2.2d depict the FO model objective function values for different values of T_1 and T_2, respectively. Total system cost is minimized when $T_1 = 6$ (Figure 2.2b) and $T_2 = 11$ days (Figure 2.2d). In other words, for a clinic of this size, experiencing this

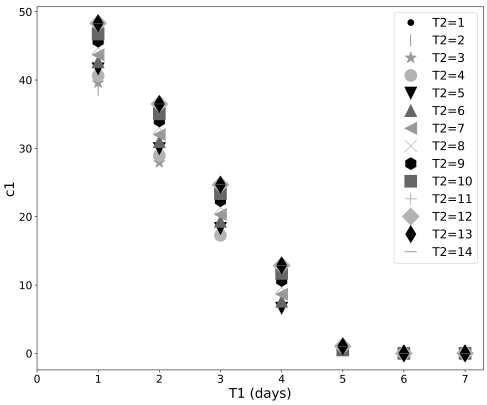

(a) Unit wait time cost for priority 1 patients as a function of T_1 for different values of T_2

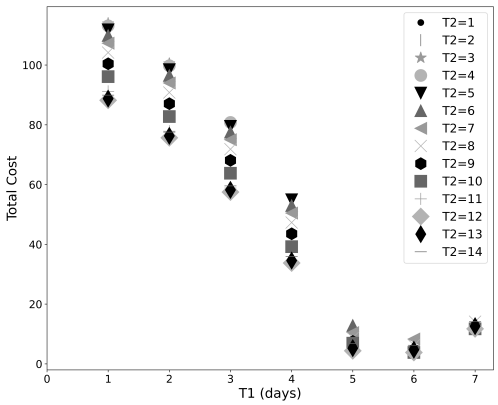

(b) Total cost as a function of T_1 for different values of T_2

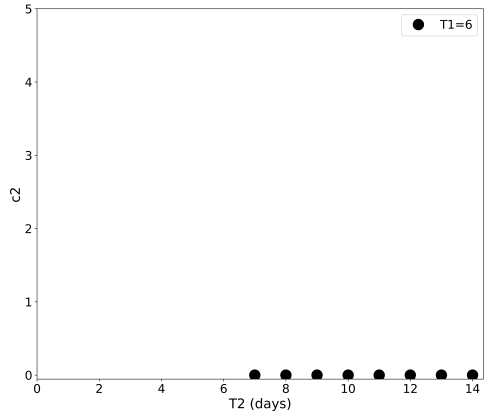

(c) Unit wait time cost for priority 2 patients as a function of T_2 for $T_1 = 6$

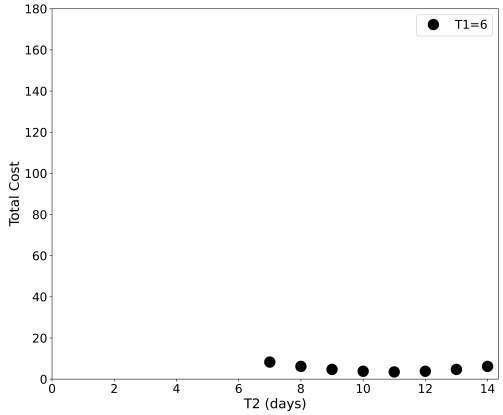

(d) Total cost as a function of T_2 for $T_1 = 6$

Figure 2.2: Unit Wait Time Cost Values (c_1 and c_2) and Total Cost for the Clinical Setting with Two Patient Classes (base-case)

demand and utilizing the DMB policy, there is no positive choice of c_1 and c_2 that could justify increasing T_1 and T_2 beyond 6 and 11 days, respectively. Put another way, one would have to value patient waiting time at essentially zero in order to justify targets beyond 6 and 11 for this setting.

For the two-patient class setting with the first available slot scheduling policy, we first need to determine booking limits, which are defined as the number of days beyond the wait time target when the cost of waiting becomes higher than that of overtime. This is typically done by dividing the unit overtime cost by the unit wait time cost associated with each priority class. In our case, however, the unit wait time costs are not known in advance. For this reason, we decided to choose 2 and 10 days as the booking limits for priority 1 and 2 patients, respectively. These values are based on the maximum cost of waiting determined through the use of the DMB policy for the two-priority class setting (base-case), which are approximately $c_1 = \$50$ for priority 1 and $c_2 = \$10$ for priority 2 patients when the overtime cost is \$100.

The total cost attains its minimum value at $T_1 = 7$ days (Figure 2.3a), and it remains unchanged with an increase in the values of T_2 (Figure 2.3b). Because T_2 was not statistically significant it does not have an impact on the total cost. In other words, using a first available slot policy increases the wait time target of the highest priority class and subsumes the priority 2 class into the first. This is intuitively reasonable as a first available slot policy essentially ignores the priority and thus needs a larger window to distribute the single class of patients. In a scenario with an imbalanced system, the total cost is minimized when the value of T_1 is 3 days and remains steady thereafter (Figure 2.3c). For the second class, $T_2 > T_1 = 3$ will have the same cost regardless of the value of T_2 (Figure 2.3d). Thus, in a system with insufficient capacity, targets are actually shortened as overtime is inevitable. Idle time is unlikely leading to less need to pre-fill the schedule and thus allowing the clinic to function with shorter wait time targets without increasing overtime costs but, obviously, at a higher overtime cost than one sees in the (balanced) base scenario.

2.4.2 Three Priority Class Setting

Figure 2.4 presents the total system cost for different choices for the wait time targets. Each of these figures shows diminishing cost savings as the wait time targets are increased. The total cost decreases over the entire range of T_2 and T_3 as the value of T_1 increases (Figure

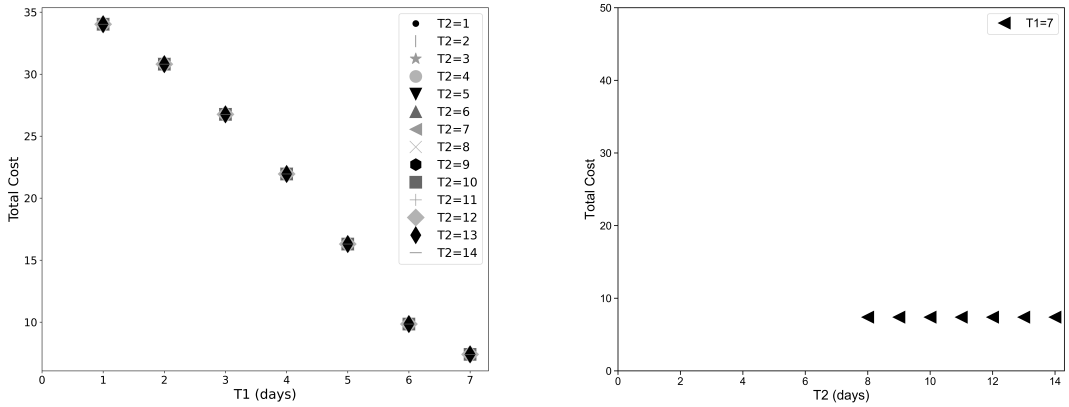

(a) Total cost as a function of T_1 for different values of T_1 (First Available Slot Policy)

(b) Total cost as a function of T_2 for $T_1 = 7$ days (First Available Slot Policy)

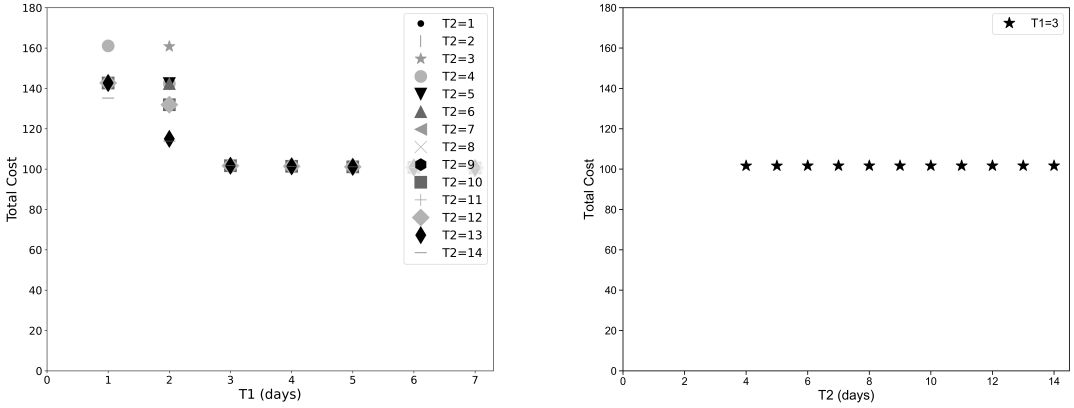

(c) Total cost as a function of T_1 for different values of T_2(Imbalanced System)

(d) Total cost as a function of T_2 for $T_1 = 3$ days (Imbalanced System)

Figure 2.3: Total Cost for the Clinical Setting with Two Patient Class Variants (Imbalanced System and First Available Slot Policy Cases)

2.4a) attaining its lowest cost value by $T_1 = 5$ days. In addition, for a fixed T_1, total cost reaches the minimum value around $T_2 = 10$ days, suggesting that no further cost savings can be obtained by setting this target beyond this point (Figure 2.4b). Finally, given $T_1 = 5$ and $T_2 = 10$ days, total cost versus T_3 decreases gradually and attains the marginal cost savings of less than 2% at $T_3 = 18$ days (Figure 2.4c).

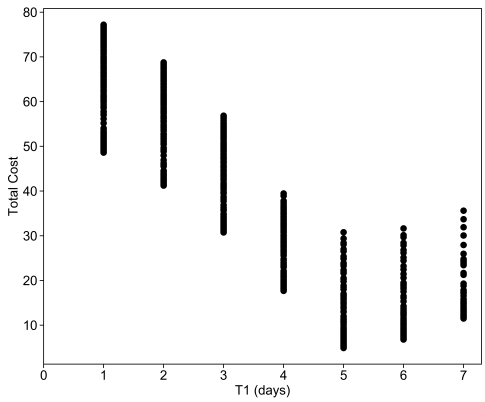

(a) Total Cost versus T_1 for different values of T_2 and T_3

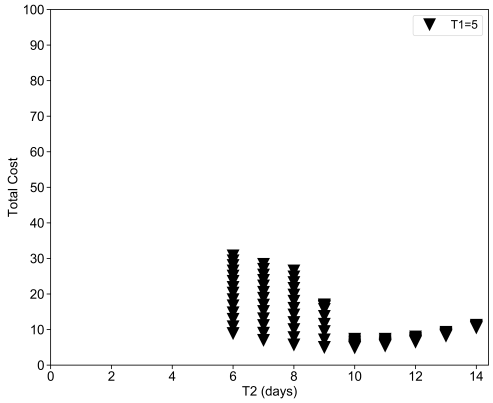

(b) Total Cost versus T_2 for different values of T_3 and $T_1 = 5$ days

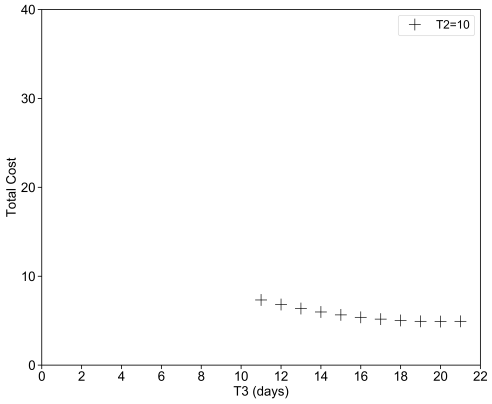

(c) Total Cost versus T_3 for $T_1 = 5$ and $T_2 = 10$ days

Figure 2.4: Total Cost versus the different wait time target for three patient classes

2.4.3 Four Priority Classes: The Case of A Rheumatology Clinic

For the Reumatology clinic, the total cost is minimized when $T_1 = 2$ and $T_2 \geq 4$ weeks (Figure 2.5a and 2.5b). On the other hand, when $T_1 = 2$ and $T_2 \geq 4$ weeks additional

(i.e., marginal) cost savings obtained beyond $T_3 = 5$ weeks is less than 2% (Figure 2.5c). Finally, analyzing the impact of varying T_4 suggests that there is little value to including this 4th priority class (let alone the 5th that was in the original data) and that patients of this class ought to be categorized as priority 3. Thus, our method would suggest that the rheumatology clinic subsume priority classes 3,4 and 5 into one and use wait time targets of 2, 4 and 5 respectively for the three remaining priority classes. This is in stark contrast to the current 5 targets of 2, 9, 13, 26, and 52. Clearly, current practice forces significantly longer waits than is justifiable from a managerial perspective.

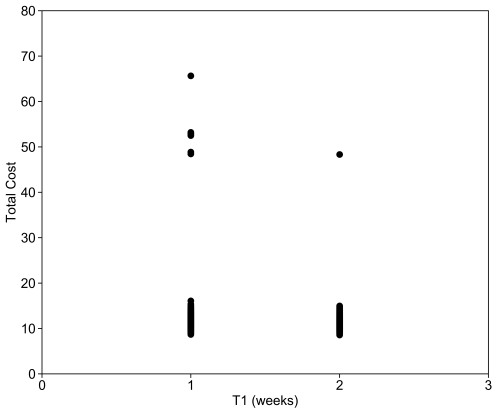

(a) Total Cost versus T_1 for different values of T_2, T_3 and T_4

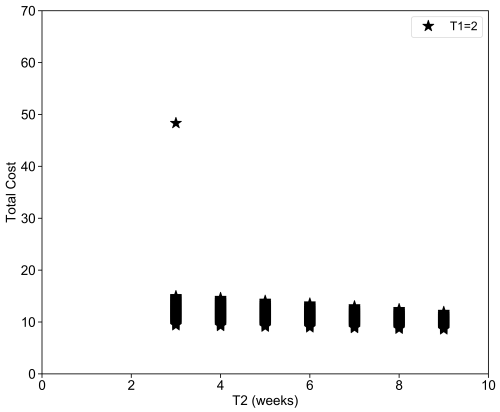

(b) Total Cost versus T_2 for different values of T_3 and T_4, and for $T_1 = 2$ weeks

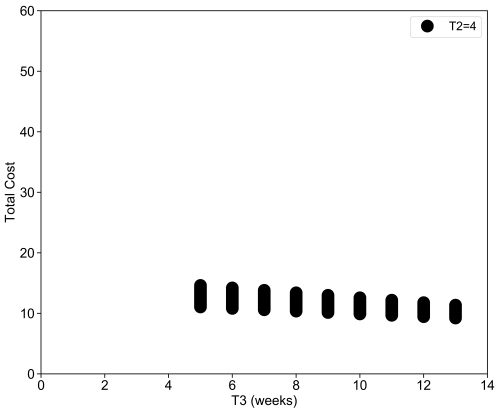

(c) Total Cost versus T_3 for different values of T_4 and for $T_1 = 2$ and $T_2 = 4$ weeks

Figure 2.5: Total Cost versus the different wait time target for a case study consisting of four priority classes

2.5 Discussion and Conclusion

Health care institutions face numerous challenges in providing care to patients within maximum recommended wait times. The typical approach to deal with this problem is to lengthen the wait time target for the lowest priority patients, an approach that provides only temporary relief. In this chapter, we described a rigorous approach to set the wait time targets in a systematic fashion that both respects clinically acceptable wait time targets and ensures that wait times are not unnecessarily lengthened. our proposed approach provides a means for determining the minimum values for the targets that achieve the near maximal reduction possible in overtime utilization and hence, a cost-savings opportunity.

Our approach can be summarized as follows. First, determine a scheduling policy. In the chapter, we discuss the performance of several advance scheduling policies found in the literature and settle on one. Other users may choose an alternative policy. Second, we simulate the scheduling policy given a known demand distribution and capacity and fit regression and neural network models to approximate the convex functions that describe the relationship between mean wait times, mean overtime and the wait time targets in a multi-priority patient setting. Finally, using the parameterized cost functions, we formulate forward and inverse mathematical problems to determine when the unit wait time cost drops to zero as wait time target values increase. This represents the minimum values for the targets that achieve the maximal reduction possible in overtime utilization.

Through numerical analyses, we show that typically the optimal wait time targets are quite low. For instance, in the case of a system with two priority classes (base-case) and a capacity of 10 appointment slots per day and that is employing the DMB scheduling policy, we find that there is no additional benefit in setting the targets beyond $T_1 = 6$ and $T_2 = 11$ days, respectively. Similarly, for a system involving three priority classes and a capacity of 10, setting T_1, T_2 and T_3 beyond 5, 10, and 18 days, respectively, does not yield additional benefits to the system and thus the additional wait cannot be justified. It is worth noting that our methodology to determine the optimal targets is general but the results presented in this chapter are case specific as it takes into account the capacity, demand and the scheduling policy employed. Generalizable results would defeat the purpose of providing targets that are achievable within a given setting. For example, it is insightful to note that classical regression models performed poorly for the rheumatology case study

- most likely due to the fact that the lower priority classes constitute a high proportion of the total demand whereas in the illustrative examples it was the highest priority class that accounted for most of the appointment requests. In addition, it is also important to note that first priority class with highest proportion of total demand results in longer optimal wait time target than the first priority class with the smallest proportion of total demand. This highlights the fact that the setting of targets needs to be case specific as what is reasonably achievable will depend on the capacity, the demand stream and the scheduling policy employed. Unlike the consensus-based approach that only takes into account medical necessity, our approach ensures that targets are reached while also ensuring that patient wait times are as low as possible. Our approach does not ignore the size or the capacity of clinic, however, as a future extension of this work, it would be worthwhile to evaluate the robustness of the results with different clinic sizes.

One question we have not particularly addressed in this work is that "what happens if the targets determined through the consensus approach in practice are shorter than the optimal targets?". Most likely the inverse optimization approach will yield optimal targets that are equal to the clinical maximums. This result will suggest that there is a cost-savings opportunity should the clinic choose to extend the clinical maximum targets.

In current practice, the overall rationale for the choice of wait time targets is not clearly laid out. With the current wait time targets, patients may wait longer for no practical benefit to hospitals or clinics in terms of resource management. The Table 2.1 demonstrates several examples of services with long clinical wait time targets for lower priority patients set through the consensus approach. As we demonstrated using illustrative examples and clinical case study, by utilizing the modelling approach in this chapter, we contend that many of the current wait time targets could be reduced significantly without requiring additional resources. (Of course, it is certainly true that an initial outlay of resources is required to reduce the wait list to the point where the new targets can be maintained.) Finally, our work highlights the important role that the scheduling policy and other operational factors such as demand and capacity play in determining what wait time targets that are readily achievable.

We tested two patient scheduling policies as part of the first step of the wait time target setting process, namely the DMB policy and a Myopic policy. Myopic policy has different variants. In this chapter, we used the First Available Slot policy with booking limits variant.

However, there are other patient scheduling policies including but not limited to the PPQ and Protection-Level policies. As a future extension of this work, it would be worthwhile to explore how different policies affect the determination of optimal wait time targets. Potential challenges might arise from the approximation of functions and implementation of mathematical programs with different functions.

The most notable benefit that can be obtained from the implementation of this approach in clinical practice is the reduction of wait times in a systematic way ensuring timely patient access to care. As emphasized in the chapter, this can be achieved through the use of good scheduling policies and accounting for demand and capacity at the specific clinical setting. To this end, the DMB policy is an easy to implement scheduling heuristic with provably better system and service level performance than other known policies in the literature. Using a DMB policy manually in practice can still be difficult and resource intensive due to a need for extensive time that booking clerks will spend for closely examining the appointment schedule and determining the first day with the minimum number of bookings in the booking horizon. An appointment scheduling software integrated with the patient information system can be useful to implement the DMB policy and ensure efficient scheduling operations. To further automate the process, an integrated decision-support tool with the optimization engine can be developed to assist with the determination of the wait time targets based on the priority classes and the associated clinical maximums, demand rates and capacity parameters for a particular clinic.

2.6 Bibliography

[1] Henry Ahn, Hans Kreder, Nizar Mahomed, Dorcas Beaton, and James G Wright. Empirically derived maximal acceptable wait time for surgery to treat adolescent idiopathic scoliosis. *CMAJ*, 183(9):E565–E570, 2011.

[2] Ravindra K Ahuja and James B Orlin. Inverse optimization. *Operations Research*, 49 (5):771–783, 2001.

[3] Dominique Ansell, James AG Crispo, Benjamin Simard, and Lise M Bjerre. Interventions to reduce wait times for primary care appointments: a systematic review. *BMC health services research*, 17(1):295, 2017.

[4] Davood Astaraky and Jonathan Patrick. A simulation based approximate dynamic programming approach to multi-class, multi-resource surgical scheduling. *European Journal of Operational Research*, 245(1):309–319, 2015.

[5] Luciana Ballini, Antonella Negro, Susanna Maltoni, Luca Vignatelli, Gerd Flodgren, Iveta Simera, Jane Holmes, and Roberto Grilli. Interventions to reduce waiting times for elective procedures. *The Cochrane Library*, 2015.

[6] Bacchus Barua. Waiting Your Turn: Wait time for Health Care in Canada, (2017 edition), 2017.

[7] Bacchus Barua and David Jacques. The Private Cost of Public Queues for Medically Necessary Care, 2019 edition, 2019.

[8] Dimitris Bertsimas, Vishal Gupta, and Ioannis Ch Paschalidis. Inverse optimization: A new perspective on the black-litterman model. *Operations research*, 60(6):1389–1403, 2012.

[9] John R Birge, Ali Hortaçsu, and J Michael Pavlin. Inverse optimization for the recovery of market structure from market outcomes: An application to the miso electricity market. *Operations Research*, 65(4):837–855, 2017.

[10] Timothy CY Chan, Tim Craig, Taewoo Lee, and Michael B Sharpe. Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research*, 62(3):680–695, 2014.

[11] Karen Davis, Kristof Stremikis, David Squires, and Cathy Schoen. Mirror, mirror on the wall. *How the performance of the US Health care system compares internationally. New York: CommonWealth Fund*, 2014.

[12] Sarah Derrett, Charlotte Paul, and Jenny M Morris. Waiting for elective surgery: effects on health-related quality of life. *International Journal for Quality in Health Care*, 11(1):47–57, 1999.

[13] Zeynep Erkin, Matthew D Bailey, Lisa M Maillart, Andrew J Schaefer, and Mark S Roberts. Eliciting patients' revealed preferences: an inverse markov decision process approach. *Decision Analysis*, 7(4):358–365, 2010.

[14] Na Geng and Xiaolan Xie. Optimal dynamic outpatient scheduling for a diagnostic facility with two waiting time targets. *IEEE Transactions on Automatic Control*, 61 (12):3725–3739, 2016.

[15] Anthony J Harrison and John Appleby. English NHS waiting times: what next? *Journal of the Royal Society of Medicine*, 102(7):260–264, 2009. PMID: 19605856.

[16] Anthony J Harrison and John Appleby. Optimising waiting: a view from the English National Health Service. *Health Economics, Policy and Law*, 5(4):397–409, 2010.

[17] Junfei Huang, Boaz Carmeli, and Avishai Mandelbaum. Control of patient flow in emergency departments, or multiclass queues with deadlines and feedback. *Operations Research*, 63(4):892–908, 2015.

[18] Garud Iyengar and Wanmo Kang. Inverse conic programming with applications. *Operations Research Letters*, 33(3):319–330, 2005.

[19] Pooyan Kazemian, Mustafa Y Sir, Mark P Van Oyen, Jenna K Lovely, David W Larson, and Kalyan S Pasupathy. Coordinating clinic and surgery appointments to meet access service levels for elective surgery. *Journal of biomedical informatics*, 66:105–115, 2017.

[20] Arezou Keshavarz, Yang Wang, and Stephen Boyd. Imputing a convex objective function. In *2011 IEEE International Symposium on Intelligent Control*, pages 613–619. IEEE, 2011.

[21] Siqiao Li, Na Geng, and Xiaolan Xie. Radiation queue: Meeting patient waiting time targets. *IEEE Robotics & Automation Magazine*, 22(2):51–63, 2015.

[22] Nan Liu. Optimal choice for appointment scheduling window under patient no-show behavior. *Production and Operations Management*, 25(1):128–142, 2016.

[23] Siciliani Luigi, Borowitz Michael, and Moran Valerie. *OECD Health Policy Studies Waiting Time Policies in the Health Sector What Works?: What Works?* Oecd Publishing, 2013.

[24] Xiang Ma, Antoine Sauré, Martin L Puterman, Marianne Taylor, and Scott Tyldesley. Capacity planning and appointment scheduling for new patient oncology consults. *Health care management science*, 19(4):347–361, 2016.

[25] Jeffrey L Mahon, Robert B Bourne, Cecil H Rorabeck, David H Feeny, Larry Stitt, and Susan Webster-Bogaert. Health-related quality of life and mobility of patients awaiting elective total hip arthroplasty: a prospective study. *Canadian Medical Association Journal*, 167(10):1115–1121, 2002.

[26] Ugenthiri Naiker, Gerry FitzGerald, Joel M Dulhunty, and Michael Rosemann. Time to wait: a systematic review of strategies that affect out-patient waiting times. *Australian Health Review*, 42(3):286–293, 2018.

[27] William G Paterson, William T Depew, Pierre Paré, Denis Petrunia, Connie Switzer, Sander J Veldhuyzen van Zanten, Sandra Daniels, Canadian Association of Gastroenterology Wait Time Consensus Group, et al. Canadian consensus on medically acceptable wait times for digestive health care. *Canadian Journal of Gastroenterology and Hepatology*, 20(6):411–423, 2006.

[28] Jonathan Patrick, Martin L Puterman, and Maurice Queyranne. Dynamic multipriority patient scheduling for a diagnostic resource. *Operations research*, 56(6):1507–1525, 2008.

[29] Claudia Sanmartin, François Gendron, Jean-Marie Berthelot, and Kellie Murphy. Access to health care services in canada, 2003. *Ottawa: Statistics Canada*, 2004.

[30] Antoine Sauré, Jonathan Patrick, and Martin L Puterman. Simulation-based approximate policy iteration with generalized logistic functions. *INFORMS Journal on Computing*, 27(3):579–595, 2015.

[31] Antoine Sauré, Jonathan Patrick, Scott Tyldesley, and Martin L Puterman. Dynamic multi-appointment patient scheduling for radiation therapy. *European Journal of Operational Research*, 223(2):573–584, 2012.

[32] Joseph Schaafsma. Are there better ways to determine wait times? *Cmaj*, 174(11): 1551–1552, 2006.

[33] Jianzhong Zhang and Chengxian Xu. Inverse optimization for linearly constrained convex separable programming problems. *European Journal of Operational Research*, 200(3):671–679, 2010.

[34] Liping Zhou, Na Geng, Zhibin Jiang, and Xiuxian Wang. Multi-objective capacity allocation of hospital wards combining revenue and equity. *Omega*, 81:220–233, 2018.

3 Dynamic Advance Patient Scheduling with Follow-up Appointments

In this chapter, we consider an outpatient clinic setting where patients require multiple appointments with a care provider. In this setting, the first visit is usually defined as a consult and possible subsequent visits are known as follow-up appointments. The goal here is to dynamically allocate the outpatient clinic capacity between consult and follow-up appointments in a cost-efficient manner with the help of a good and effective scheduling policy.

3.1 Introduction and Motivation

Healthcare spending in Canada continues to trend upward reaching 11.6% of the gross domestic product (GDP) in 2019. A rise in healthcare expenditure can potentially limit the ability of governments to provide other services such as education and social assistance programs. The Canadian Medical Association (CMA), in a recent report, indicated that 59% of Canadians agree that this trend is unsustainable and 91% of these people believe that increasing the efficiency of clinical settings is the best way to slow the growing costs [3]. We believe there is a significant room for healthcare providers to improve efficiency, quality and access to the medical services.

Scheduling in healthcare is complex mainly due to the random nature of patient arrivals and service times, the existence of multiple and inter-connected episodes of care, and limited constrained healthcare resources. It is also a critical part of healthcare operations and has a direct impact on timely access and delivery of care.

There are two main streams of patient scheduling problems. Advance scheduling (i.e., inter-day scheduling) focuses on waiting costs incurred until the day of service whereas appointment scheduling (i.e., intra-day scheduling) focuses on costs incurred on the day of service (e.g., patient waiting, clinic overtime and idle time costs). There has been significant

research in both streams of patient scheduling problems [22, 7, 31, 17, 15, 34, 18, 38, 4]. Most of the prior literature in this area has focused on scheduling of a single-appointment for each patient where each visit to the clinic is treated independently. In this study, we consider an advance scheduling setting in which a patient may visit the provider multiple times. The scheduling is even more difficult in this setting because of the challenging task of prioritizing and allocating a fixed service capacity between new and returning patients of different types. In general, patients require multiple visits with some pre-defined frequency, usually specified by medical guidelines. In other words,they require an initial consult visit followed by potentially one or more follow-up visits.

Follow-up practices differ, both in the timing and the frequency of visits depending on the type of care and physicians. Inefficient follow-up practices can result in longer waits for new and existing patients which in turn can negatively impact patient outcomes [8, 14, 24]. The question of how often a patient must be seen is both a clinical and an operational question since it is directly associated not only with the clinical needs or patient outcomes but also with the care provider's workload, staffing levels and hence healthcare costs. In general, the length of follow-up intervals is an under explored area both in the medical and the operations research literature. There are several clinical and management research studies that investigate the appropriate length of follow-up intervals in various settings [39, 2, 6, 25, 27]. Due to several clinical challenges there is no clear consensus regarding follow-up visit guidelines for many chronic or curable health conditions [37, 28, 35, 11, 23].

There are several outpatient health services delivered through a series of appointments. Chemotherapy drugs and many other emerging cancer drugs are usually administered via an intravenous infusion in multiple visits following the short intervals. For example, adult patients with Hodgkin lymphoma receive brentuximab vedotin every three weeks for up to 16 cycles. Radiotherapy require that cancer patients to undergo radiation sessions on consecutive days until the necessary radiation dose is given to the targeted tumour. Mental health outpatient services are also delivered via multiple visits to psychiatrists or therapists. Over the years, mental health outpatient clinics have experienced significant wait times resulting in cancelled follow-up visits for many patients. The Choice and Partnership Approach (CAPA) is a service model adopted by mental health service providers with the aim to ensure effective use of the capacity for outpatient services. According to this model, young and adult patients need a choice (i.e., consult) appointment and on average

7.5 partnership (i.e., follow-up) appointments, where two thirds of the patients who have a choice appointment will require partnership appointments. Moreover, diabetes, dialysis, physiotherapy and dental/orthodontic treatment settings require care to be delivered through several appointments.

This research aims to partly fill the existing knowledge gap as to the impact of follow-up guidelines on patient wait times and resource capacity utilization. The objective of this research is to help determine consult and follow-up scheduling policies that take into account a resource's availability at multiple points in time. Specifically, we seek to answer the following research questions: How should patients be scheduled both for initial consult and follow-up appointments? Alternatively, given a pre-specified frequency and number of follow-up visits as per existing guidelines, what is a good scheduling policy that could allow the care providers to meet wait time targets for new consults while keeping follow-up appointments within target time windows?

We believe the research described in this chapter here has the potential to inform more efficient patient follow-up practices for ambulatory care services and hence lead to improved patient outcomes by providing patients with the right care at the right time. The rest of this chapter is structured as follows. In Section 3.2, we briefly describe the related literature on multi-appointment patient scheduling. In Section 3.3, we describe Markov decision process (MDP) formulation for patient scheduling with multiple appointments. We note that due to the model intractability and computational issues, the proposed MDP model and its LP equivalent cannot be solved for large instances. Then, we solve the model approximately for small problem settings and identify the characteristics of the approximate optimal policy (AOP). Based on the AOP decisions, we define a heuristic that can be applied to a large-scale practical settings. Finally, we provide a numerical analysis using data from a local hospital and compare the performance of the proposed heuristic policy against that of other policies via simulation. In the last section, we describe next steps and future work.

3.2 Related Literature

Most studies model advance patient appointment scheduling as a stochastic optimization problem with a trade-off between capacity utilization and waiting times and other service level measures. There are several studies in the literature that examined the advanced

scheduling of single visit for a single resource. Patrick et al. (2008) developed an MDP model to schedule patient appointments for a diagnostic CT scan service and derived a set of booking rules now called the PPQ policy [31]. According to the PPQ policy, high priority patients are booked as soon as possible whereas lower priority patients are booked on day 1, then on the last day within their maximum recommended wait, and then booked down to day 2 in an order of the urgency. Overtime capacity is used when there is insufficient capacity within the target. Sauré et al. (2015) compared via simulation the performance of several advance scheduling policies including the PPQ policy and demonstrated that The Day with the Minimum number of Bookings (DMB) policy had a superior performance compared to PPQ and Myopic policies [33]. Several others such as Green et al. (2006), Erdeyli and Topaloglu (2009) and Dobson et al. (2011) also framed the advance scheduling problem as a dynamic optimization problem with a trade-off between capacity utilization and waiting times or between revenue and service levels [21, 16, 15]. Green et al. (2008) and Liu et al. (2013) addressed patient scheduling as a queuing problem in a primary care setting [20, 29]. Several other papers in the literature consider advance scheduling problems involving multiple healthcare resources [18, 38, 4].

The number of studies involving multiple appointments is rather limited. Using an MDP formulation, Diamant et al. (2018) proposed a scheduling policy when the decisions for multiple assessments are made on the day of the appointment [12]. While reducing the negative effect of no-shows, this policy provides the flexibility to schedule patients' next assessment on the day of the booked appointment. In this vein, Marynissen and Demeulemeester (2018) provide a comprehensive review of the literature on multi-appointment patient scheduling problem where patients need to undergo multiple diagnostic tests, specialist consultations, and surgeries at multiple units of a clinical setting [30]. Leefthink et al. (2020) provide a comprehensive review on multi-disciplinary appointment scheduling problems which include multiple interrelated appointments per patient [26].

In a related paper, Sauré et al. (2012) considered a radiotherapy capacity planning setting with 18 different types of treatment in which patients need to undergo a certain number of treatment sessions on consecutive days [34]. Using dynamic programming and value function approximation techniques, they were able to identify cost-effective guidelines for allocating available treatment capacity to patients requiring a number of sessions and a number of slots per session. Gocgun and Puterman (2014) considered a chemotherapy

scheduling problem in which treatment sessions need to be booked at predefined time intervals as mandated by the treatment protocols for cancer type [19]. In another related paper, Ding et al. (2020) studied follow-up appointments using a re-entrant queue, and proposed a policy to designate returning customers as high-priority ones, placing them at the head of the queue [13].

The most related work to ours is a recent paper by Yu et al. (2020) [40]. They studied a physical therapy clinic setting with multiple visits and formulated an average reward MDP model to maximize the long-term profit (revenue - cost) associated with each patient. Our work shares a few common characteristics with theirs. First, we also allow both deterministic and stochastic number of visits. Second, we also note that the number of visits can follow a Geometric distribution. However, our study bears significant differences compared to theirs. First, they solve an average reward MDP model by making the uni-chain Markov chain structure assumption whereas we solve a discounted MDP with discount factor close to 1 as a way of solving the average-cost problems exactly or approximately. In practice, the irreducibility or reachability property of Markov chains depends on the booking policies being used. The advantage of solving a discounted cost MDP is that the theory holds more generally and there is no need to make strong assumptions on uni-chain policies. Also, numerical methods demonstrate better convergence for discounted cost/reward MDPs. Second, their model does not support wait time windows for consult visits greater than the wait time window for the follow-up, whereas our model does not constrain the wait time targets for consult or follow-up visits. Third, while they consider a small setting with two heterogeneous patient classes based on the number of visits and frequency, we allow for multiple patient priority classes defined based on clinical urgency, wait time targets, and the follow-up visit interval. Finally, the MDP policy we derive is heuristic-based which makes its implementation easier in practical settings unlike the Index Policy (IP) they propose which requires more computational effort.

3.3 Markov Decision Process Formulation

This section provides the detailed formulation of a discounted infinite-horizon MDP model proposed for this problem.

3.3.1 Decision Epoch & Booking Horizon

We consider a system that has a regular-hour capacity of Cr fixed-length appointment slots and an overtime capacity of Co fixed-length appointment slots. For modelling convenience, we assume weekly scheduling decisions to be able to handle long booking horizons. Each week a booking agent receives requests from new patients for appointments with varying degrees of clinical urgency and observes the booking schedule with already booked patients who might require follow-up visits. At each point, the booking agent can book first visits and the follow-up visits at most N weeks in advance where N is larger than or equal to the wait time target of the lowest priority class.

3.3.2 State Space

The state space $(\vec{x}, \vec{u}, \vec{y}) = (x_{111}, \ldots, x_{IKN}, u_{111}, \ldots, u_{IKN}, y_{11}, \ldots, y_{IK})$ is composed of a vector $\vec{x}$ where $x_{ikn} \in Z_0^+$ is the number of priority i patients already booked in week n for their kth visit using regular-hour capacity, a vector $\vec{u}$ where $u_{ikn} \in Z_0^+$ is the number of priority i patients already booked using in week n for their kth visit using overtime, and a vector $\vec{y}$ where $y_{ik} \in Z_0^+$ is the number of priority i patients waiting to be booked for their kth visit (where $k = 1$ indicates a first visit or consult and $k > 1$ is a subsequent visit or follow-up). We assume that the state variables are bounded above by the corresponding capacities or a maximum demand to ensure a finite state space.

3.3.3 Action Space

At the beginning of each week, the booking agent must decide how many consults and follow-up visits to book into the next N weeks. The booking action $(\vec{a}, \vec{z}) \in (a_{111}, \ldots, a_{IKN}, z_{111}, \ldots, z_{IKN})$ is composed of two elements. The action variable $a_{ikn} \in Z_0^+$ defines the number of priority i patients for whom appointment visit k is booked in week n using regular-hour capacity and $z_{ikn} \in Z_0^+$ specifies the number of priority i patients for whom an appointment visit k is booked in week n using overtime. We assume that all service requests receive an appointment either through regular-hour or overtime capacity, and hence there is no demand postponed to the next week. Actions must satisfy the following constraints.

$$\sum_{i,k} x_{ikn} + \sum_{i,k} a_{ikn} \leq C_r \qquad n \leq N \tag{3.1}$$

$$\sum_{n} a_{ikn} + \sum_{n} z_{ikn} = y_{ik} \qquad i \leq I, k \leq K \tag{3.2}$$

$$\sum_{i,k} u_{ikn} + \sum_{i,k} z_{ikn} \leq C_o \qquad n \leq N \tag{3.3}$$

Constraint 3.1 ensures that the number of regular-hour bookings does not exceed the regular-hour capacity. Both consult and follow-up patient appointments must be booked through either regular-hour or overtime as per Equation 3.2. Finally, constraint 3.3 ensures that the number of overtime bookings does not exceed the overtime capacity.

3.3.4 Transition Probabilities

We consider two sources of uncertainty, both of them represented with the random variable Y_{ik}, the number of requests for a consult visit ($k = 1$), and the number of requests for follow-up visits ($k > 1$). The number of new consult requests from priority i patients (Y_{i1}) follows a Poisson distribution with mean m_i. The number of patients already booked requiring a subsequent follow-up appointment (i.e., Y_{i2+}) follows a Binomial distribution with parameters (T_{ik}, p_{ik}), where $T_{ik} = x_{ik1} + a_{ik1} + u_{ki1} + z_{ik1}$ is the total number of patients having their kth visit in week 1 at the end of the current week (post-decision state variable) and p_{ik} is the probability that a priority i patient requires an additional visit (i.e., follow-up) after the kth visit.

$$p(\vec{s'}|\vec{s},\vec{a}) = \begin{cases} \prod_{i=1}^{I} \prod_{k=1}^{K} \mathbb{P}(Y_{ik} = y_{ik}), & \text{if } \vec{s'} \text{ satisfies the contraints (3.4) - (3.7)} \\ 0 & \text{otherwise.} \end{cases}$$

$$x'_{ikn} = x_{ikn+1} + a_{ikn+1} \qquad i \leq I, k \leq K, n < N \tag{3.4}$$

$$u'_{ikn} = u_{kin+1} + z_{ikn+1} \qquad i \leq I, k \leq K, n < N \tag{3.5}$$

$$x'_{ikN} = u'_{ikN} = 0 \qquad i \leq I, k \leq K \tag{3.6}$$

$$y'_{ik} = y_{ik} \qquad i \leq I, k \leq K \tag{3.7}$$

Since we consider a rolling booking horizon, at the beginning of each week, the number of appointments booked in the Nth week is zero.

3.3.5 Immediate Costs

The cost of taking action $\vec{a}$ in state $\vec{s}$ arises from the following sources: the cost associated with booking consult visits beyond the recommended wait time targets, the cost associated with the booking follow-up visits outside of the recommended wait time window, and the cost associated with overtime capacity usage.

$$c\left(\vec{s}, \vec{a}\right) = \sum_{i,k,n} c_{ikn} a_{ikn} + \sum_{i,k,n} (c_{ikn} + h) z_{ikn} \tag{3.8}$$

There is a weekly penalty of $g(i)$ for booking a consult appointment ($k = 1$) for priority i patients beyond the wait time target, where the wait time target is denoted by $T_c(i)$.

$$c_{ikn} = \begin{cases} \left(n - T_c(i)\right) g(i), & \text{if } n \geq T_c(i) \quad k = 1; \\ 0. & \text{otherwise.} \end{cases} \tag{3.9}$$

There are also weekly penalties of c_l and c_e for booking a follow-up appointment ($k > 1$) outside of the recommended wait time window defined by tolerances $lb(i)$ and $ub(i)$ where $T_f(i)$ is the ideal wait time for follow-up visit.

$$c_{ikn} = c_l \times max\left[n - T_f(i) - ub(i), 0\right] + c_e \times max\left[T_f(i) - lb(i) - n, 0\right] \quad k > 1 \tag{3.10}$$

Finally, there is a fixed overtime cost of h per appointment slot.

3.3.6 Optimality Equations

The value function of the MDP model represents the minimum expected total discounted cost over an infinite time horizon for each state and satisfies the following Bellman's optimality equations:

$$v\left(\vec{s}\right) = \min_{\vec{a} \in A_{\vec{s}}} \left\{ c\left(\vec{s}, \vec{a}\right) + \gamma \sum_{\vec{s'} \in S} p(\vec{s'}|\vec{s}, \vec{a}) v\left(\vec{s'}\right) \right\} \qquad \forall \vec{s} \in S, \tag{3.11}$$

where γ is the weekly discount factor. The dimension of the state space is $2IKN + IK$ and of the action variables is $2IKN$. Thus, any realistic values of C_r, C_o, I, N and Y_{ik} make a direct solution of the Bellman's equations impractical due to the computational intractability.

3.4 Solution Approach: The Linear Programming Approach to Approximate Dynamic Programming

Using the fundamental result presented by d'Epenoux [10], we can transform the above MDP model into an equivalent LP form as follows:

$$\max \sum_{\vec{s} \in \vec{S}} \alpha(\vec{s}) v(\vec{s}) \tag{3.12}$$

subject to

$$c(\vec{s}, \vec{a}) + \gamma \sum_{s' \vec{\in} S} p(\vec{s'}|\vec{s}, \vec{a}) v(\vec{s'}) \geq v(\vec{s}) \quad \forall \vec{s} \in S, \vec{a} \in A_{\vec{s}} \tag{3.13}$$

where $\alpha(\vec{s})$ is strictly positive. We assume that $\sum_{\vec{s} \in S} \alpha(\vec{s}) = 1$ and consider $\alpha(\vec{s})$ as probability distribution for the initial state of the system. However, the equivalent LP formulation does not avoid the issue of tractability since there is one decision variable for each state and one constraint for each feasible state-action pair. To deal with the computational intractability, we resort to the Approximate Dynamic Programming (ADP) techniques [5, 32]. More specifically, we adopt the linear programming approach to ADP [36, 9, 1] and approximate the value function $v(\vec{s})$ using the following affine approximation architecture in the

pre-decision state variables:

$$v\left(\vec{x},\vec{u},\vec{y}\right)=W_o+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}V_{ikn}x_{ikn}+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}U_{ikn}u_{ikn}+\sum_{i=1}^{I}\sum_{k=1}^{K}W_{ik}y_{ik} \qquad (3.14)$$
$$\vec{V}\geq 0,\vec{U}\geq 0,\vec{W}\geq 0,W_o\in\mathbb{R}$$

The components of the vectors $\vec{V}$ and $\vec{U}$ are directly interpretable. V_{ikn} and U_{ikn} represent the marginal expected total discounted cost of having an additional slot booked in week n with the kth visit of a priority i patient using regular-time and overtime capacity, respectively. The elements of the vector $\vec{W}$ represent the marginal expected total discounted cost of having an additional consult or follow-up visit for a priority i patient waiting to be booked. Reformulating the LP in terms of the affine approximation yields the following approximate linear program (ALP):

$$\max_{W_o\in\mathbb{R},\vec{V},\vec{U},\vec{W}}\sum_{(\vec{x},\vec{u},\vec{y})\in S}\alpha\left(\vec{x},\vec{u},\vec{y}\right)\left(W_o+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}V_{ikn}x_{ikn}+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}U_{ikn}u_{ikn}+\right.$$
$$\left.+\sum_{i=1}^{I}\sum_{k=1}^{K}W_{ik}y_{ik}\right) \qquad (3.15)$$

subject to

$$W_0+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}V_{ikn}x_{ikn}+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N}U_{ikn}u_{ikn}+\sum_{i=1}^{I}\sum_{k=1}^{K}W_iy_{ik}-$$
$$-\gamma\sum_{\vec{d}}\left[p(y'_{ik})\left(W_0+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N-1}V_{ikn}\left(x_{ikn+1}+a_{ikn+1}\right)+\right.\right.$$
$$\left.\left.+\sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n=1}^{N-1}U_{ikn}\left(u_{ikn+1}+z_{ikn+1}\right)+\sum_{i=1}^{I}\sum_{k=1}^{K}W_{ik}y'_{ik}\right)\right]\leq c\left(\vec{s},\vec{a}\right) \qquad (3.16)$$

$$\forall\left(\vec{a},\vec{z}\right)\in A_{\vec{s}},\forall\left(\vec{x},\vec{u},\vec{y}\right)\in S,\vec{V}\geq 0,\vec{U}\geq 0,\vec{W}\geq 0,\geq 0,W_o\in\mathbb{R}$$

Using the assumption that $\alpha(\vec{s})$ is a probability distribution over initial state of the system, we obtain the following ALP model:

$$\max_{W_o \in \mathbb{R}, \vec{V}, \vec{U}, \vec{W}} \left[W_0 + \sum_{i=1}^{I} \sum_{k=1}^{K} \sum_{n=1}^{N} E_\alpha[X_{ikn}] V_{ikn} + \sum_{i=1}^{I} \sum_{k=1}^{K} \sum_{n=1}^{N} E_\alpha[U_{ikn}] U_{ikn} + \sum_{i=1}^{I} \sum_{k=1}^{K} E_\alpha[Y_{ik}] W_{ik} \right] \tag{3.17}$$

subject to

$$\begin{aligned} (1-\gamma) W_0 + \sum_{i=1}^{I} \sum_{k=1}^{K} \sum_{n=1}^{N} V_{ikn} \left(x_{in} - \gamma x_{ikn+1} - \gamma a_{ikn+1} \right) + \\ + \sum_{i=1}^{I} \sum_{k=1}^{K} \sum_{n=1}^{N} U_{ikn} \left(u_{ikn} - \gamma u_{ikn+1} - \gamma z_{ikn+1} \right) + \sum_{i=1}^{I} \sum_{k=1}^{K} W_i \left(y_{ik} - \gamma E[Y_{ik}] \right) \leq c\left(\vec{s}, \vec{a}\right) \\ \forall \left(\vec{a}, \vec{z}\right) \in A, \forall \left(\vec{x}, \vec{u}, \vec{y}\right) \in S, \vec{V} \geq 0, \vec{W} \geq 0, \vec{U} \geq 0, W_o \in \mathbb{R} \end{aligned} \tag{3.18}$$

In Equation 3.17, $E_\alpha[X_{ikn}]$, $E_\alpha[U_{ikn}]$, and $E_\alpha[Y_{ik}]$ are expected values of the state variables with respect to the probability distribution α. The additional variables $x_{ikN+1}, u_{ikN+1}, a_{ikN+1}$ and z_{ikN+1} are constrained to be zero because appointments can not be booked beyond week N. The ALP has $2IKN + IK + 1$ variables and intractable number of constraints, one for each state-action pair. For this reason, we formulate its dual below.

$$\min_{\vec{X}} \sum_{(\vec{x},\vec{u},\vec{y}) \in S, (\vec{a},\vec{z}) \in A_{(\vec{x},\vec{u},\vec{y})}} X(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z}) c(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z}) \tag{3.19}$$

subject to

$$(1-\gamma) \sum_{(\vec{x},\vec{u},\vec{y}) \in S, (\vec{a},\vec{z}) \in A_{(\vec{x},\vec{u},\vec{y})}} X(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z}) = 1 \tag{3.20}$$

$$\sum_{(\vec{x},\vec{u},\vec{y})\in S,(\vec{a},\vec{z})\in A_{(\vec{x},\vec{u},\vec{y})}} X(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z})\,(x_{ikn} - \gamma x_{ikn+1} - \gamma a_{ikn+1}) \geq E_\alpha[X_{ikn}] \quad \forall i,k,n \tag{3.21}$$

$$\sum_{(\vec{x},\vec{u},\vec{y})\in S,(\vec{a},\vec{z})\in A_{(\vec{x},\vec{u},\vec{y})}} X(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z})\,(u_{ikn} - \gamma u_{ikn+1} - \gamma z_{ikn+1}) \geq E_\alpha[U_{ikn}] \quad \forall i,k,n \tag{3.22}$$

$$\sum_{(\vec{x},\vec{u},\vec{y})\in S,(\vec{a},\vec{z})\in A_{(\vec{x},\vec{u},\vec{y})}} X(\vec{x},\vec{u},\vec{y},\vec{a},\vec{z})\,(y_{ik} - \gamma E[Y_{ik}]) \geq E_\alpha[Y_{ik}] \quad \forall i,k \tag{3.23}$$

$$\vec{X} \geq 0 \tag{3.24}$$

where $E[Y_{ik}] = m_i$ for $k = 1$ and $E[Y_{ik}] = [x_{i(k-1)1} + a_{i(k-1)1} + u_{i(k-1)1} + z_{i(k-1)1}]p_{i(k-1)}$ for $k > 1$. The dual problem above has a reasonable number of constraints but an intractable number of variables, one for each state - action pair. Therefore, we solve it using column generation.

The column generation algorithm starts with a feasible state-action pair to the master problem and then uses the values of $V_{ikn}, U_{ikn}, W_{ik}, W_0$ in the sub-problem (or pricing problem) to find the most violated constraint in the primal. This procedure continues until no primal constraint is violated or a stopping criterion is satisfied (e.g., reduced cost of -0.0001). In cases where finding a single initial feasible column is not trivial, we use a phase 1 approach to find initial set of feasible columns to start the column generation algorithm.

3.4.1 Deriving a Policy from the ALP

The optimal W_0^* value does not play a role in the definition of approximate optimal policy (AOP). The optimal values of V_{ikn}^*, U_{ikn}^*, and W_{ik}^* are used to derive the AOP. To this end, we solve the following integer program obtained by inserting the optimal value function approximation (Equation 3.14) into the right hand side of Bellman's optimality equations (Equation 3.11) to determine the optimal action $(\vec{a}, \vec{z})$ given a state $(\vec{x}, \vec{u}, \vec{y})$.

$$\min_{(\vec{a},\vec{z})\in A_{(\vec{x},\vec{u},\vec{y})}} \Bigg\{ \sum_{i=1}^{I}\sum_{k=1}^{K-1}(c_{ik1}+\gamma p_{ik}W^*_{ik+1})a_{ik1} + \sum_{i=1}^{I} c_{iK1}a_{iK1} + \\ + \sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n>1}^{N}\left(c_{ikn}+\gamma V^*_{ikn-1}\right)a_{ikn} + \sum_{i=1}^{I}\sum_{k=1}^{K-1}(c_{ik1}+h+\gamma p_{ik}W^*_{ik+1})z_{ik1} + \sum_{i=1}^{I}(c_{iK1}+h)z_{iK1} + \\ + \sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{n>1}^{N}\left(c_{ikn}+h+\gamma U^*_{ikn-1}\right)z_{ikn} \Bigg\} + \text{constant} \tag{3.25}$$

The coefficients of the booking actions, denoted by A_{ikn}, are summarized below:

$$A_{ikn} = \begin{cases} c_{ik1}+\gamma p_{ik}W^*_{ik+1}, & \forall i,k \le K-1, n=1; \\ c_{ik1}, & \forall i,k=K, n=1; \\ c_{ikn}+\gamma V^*_{ikn-1}, & \forall i,k,n>1. \end{cases} \tag{3.26}$$

The overtime action coefficients, denoted by Z_{ikn} are summarized next:

$$Z_{ikn} = \begin{cases} c_{ik1}+h+\gamma p_{ik}W^*_{ik+1}, & \forall i,k \le K-1, n=1; \\ c_{ikn}+h, & \forall i,k=K, n=1; \\ c_{ikn}+h+\gamma U^*_{ikn-1}, & \forall i,k,n>1. \end{cases} \tag{3.27}$$

3.5 Results

3.5.1 Policy Insights

To derive the main characteristics of the AOP policy, we first solved a small instance of the problem consisting of two patient classes and two visits (one follow-up visit) as a base-case scenario. We then ran an extensive sensitivity analysis by changing the fixed number of visits, the conditional probability of an extra visit, the number of patient classes, wait time targets, wait time and overtime penalties, arrival rates, follow-up target wait time interval and the length of the booking horizon. The goal of this extensive sensitivity analysis was to gain insights into the main characteristics of the AOP policy, investigate the robustness of the policy while avoiding the computational effort that would be necessary to solve large

instances, and derive booking guidelines. Below, we describe the sensitivity analysis in detail.

a) *Base-case:* In the base-case scenario, we consider a hypothetical clinical setting with two patient classes and two visits (i.e., one follow-up visit) for patients where wait time targets are set at 3 and 7 weeks for priority 1 and 2 consults, respectively. The target for follow-up visits is set at 10 weeks after the consult appointment with a tolerance of two weeks. In other words, the clinic can book the follow-up appointment between week 8 and 12 of the booking horizon (i.e., 10 ± 2 weeks). We call this interval a follow-up appointment window in the remainder of this chapter. We assume 18 regular-hour appointment slots and 18 overtime appointment slots, sufficiently large number to serve the the demand not booked through regular-hour capacity. Wait time penalties for booking beyond the consult targets are \$25 and \$20 per week for priority 1 and 2, respectively. The penalty for booking outside of the follow-up appointment window is \$50 per week, and the overtime cost is \$100 per week per appointment slot. Patient arrivals follow Poisson distributions with the means of 3 and 6 requests per week for priority 1 and 2 patients, respectively. The booking horizon is 14 weeks and the discount factor is set at 0.99 as suggested earlier by Patrick et al. [31].

b) *Probabilistic follow-up visit:* We consider one consult and one follow-up visit. The need for a subsequent follow-up visit could be determined based on the combination of clinical outcomes and physician advice. Here, we assume that the follow-up visit occurs with a probability $p = 0.5$. To balance the average demand and capacity, we consider 14 regular-hour appointment slots per week. All other parameters stay the same.

c) *More follow-up visits:* We consider 3 patient visits, one consult visit followed by 2 follow-up visits, and the regular-hour capacity is 27 appointment slots per week. All other parameters stay the same.

d) *Different follow-up appointment window penalty:* The penalty for booking the follow-ups outside (both before and after the window) of the follow-up window increases from \$50 to \$100 per week per appointment slot. All the other paramaters stay the same.

e) *More patient classes:* We assume that there are 3 priority classes with consult wait time targets of 3, 7 and 10 weeks for priority 1, 2 and 3 classes, respectively. Patient arrivals follow Poisson distributions with means of 3, 6 and 4 patients requests per week, respectively. The number of regular-hour appointment slots is set at 26 per week. All the other parameters stay the same.

f) *Different overtime cost:* We vary the overtime cost as follows: \$50, \$75, and \$125 per week per appointment slot. All the other parameters remain unchanged.

g) *Different demand rates:* We consider demand rates of 8 and 6, 2 and 6, and 3 and 10 per week for priority 1 and 2 patients, respectively. The number of regular-hour appointment slots is set to match the corresponding average demand. All the other parameters stay the same.

h) *Different follow-up and booking horizon:* We consider a follow-up target of 14 weeks with an extended booking horizon of 20 weeks. All the other parameters remain the same.

i) *Different consult wait time targets:* We vary the consult wait time targets one at a time. First, we assume a wait time target of 5 weeks for the first patient class and then assume a target of 10 weeks for the second patient class. The other parameters remain unchanged.

j) *Different wait time penalties:* We change the wait time penalties from \$25 and \$20 per week for priority 1 and 2 patients, respectively, to \$20 and \$15 per week. The rest of the parameters stay the same.

k) *Different follow-up booking tolerance:* Instead of using a tolerance of -2 and $+2$ weeks, we use -1 and $+2$ weeks. All the other model parameters remain the same.

l) *Imbalanced System:* We assume that total average demand is 18 requests per week and the capacity is 17 or 16 slots per week. All the other parameters remain the same.

We implemented the column generation algorithm and the integer programming model that determines the approximate optimal actions in Java 8 with the CPLEX 12.9 as the solver. The solution of the model involved around $700 - 1,200$ columns and took between 2 and 5 minutes depending on the instance. The computer used to run all the numerical experiments in this chapter was a Intel(R) Xeon(R) PC with a 2.00 GHz processor and 32GB RAM.

Figure 3.1 summarizes the booking coefficients (i.e., A_{ikn} and Z_{ikn}) for the AOP policy in the base-case scenario. The behaviour of the AOP policy is robust with respect extensive model parameter changes in the sensitivity analysis described in scenarios (b) through (l). For the consult visits ($k = 1$), A_{ikn} is the lowest in week 1 for all priority classes. If the appointment schedule in week 1 is full, then the second lowest A_{ikn} is in the last week within the wait time target ($i.e., T_c(i)$) regardless of the number of priority classes (Figure 3.1a and 3.2a) and wait time target for each priority class (Figure 3.2b and 3.2c). Starting from

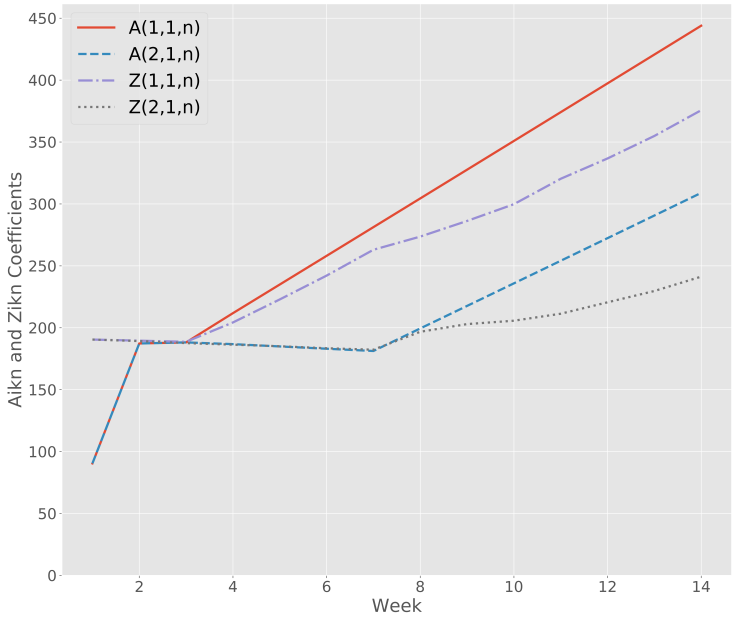

(a) consult (base-case)

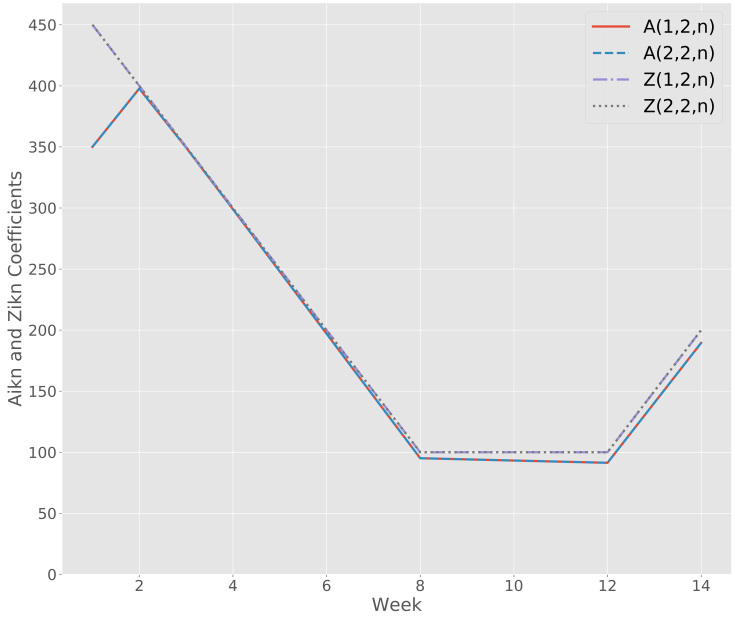

(b) follow-up (base-case)

Figure 3.1: Base-Case AOP Booking Coefficients

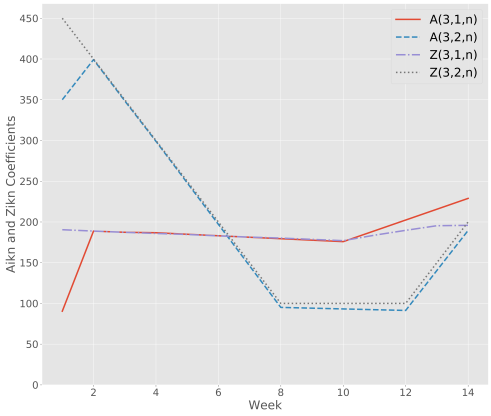

(a) Additional patient class

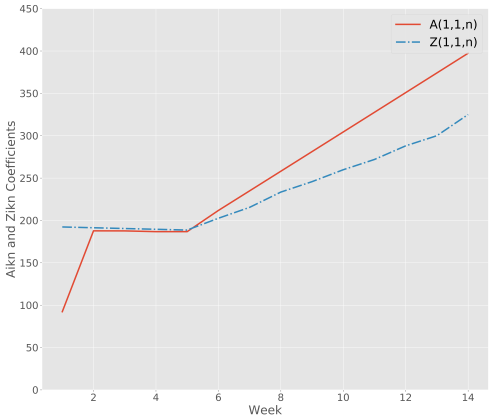

(b) $T_c(1) = 5$ weeks

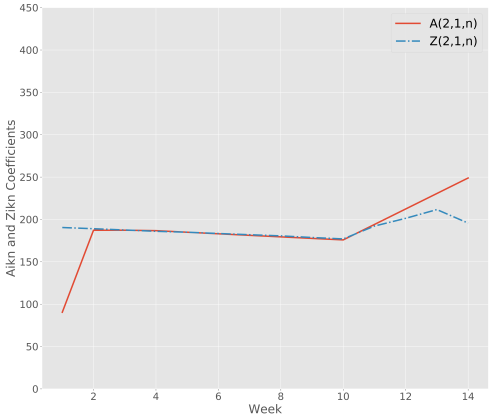

(c) $T_c(2) = 10$ weeks

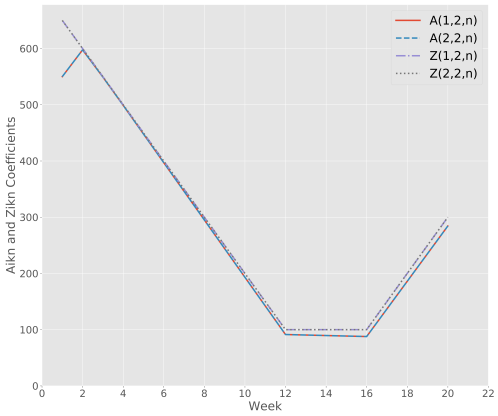

(d) $T_f(i) = 14$ weeks

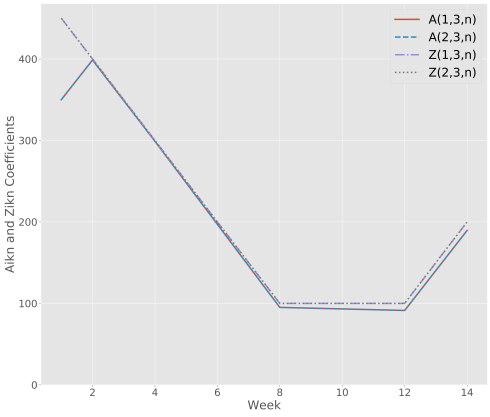

(e) Additional follow-up visit

Figure 3.2: Sensitivity Analysis to Base-Case ALP Model

this point and working down to week 1, the AOP policy suggests the use of overtime only if there is no regular-hour slot available in a particular week. In other words, the AOP policy suggests that booking all the capacity, regular-time and overtime, before moving to the next week within the wait time target. This is due to the following relationship between the different booking coefficient values $A_{ikn} \leq Z_{ikn} \leq A_{ikn-1} \leq Z_{ikn-1} \leq \ldots$. This is a consistent property of the policy for high priority patients. For low priority patients, in some instances, the AOP policy uses overtime before regular-hour capacity in the middle of the wait time target. For example, in the base-case scenario, the coefficients for the second priority class with a consult visit wait time target of $T_c(i) = 7$ weeks, $Z_{ikn} \leq A_{ikn}$ for $n = 3, 4$ and 5. A possible explanation for this observation is that the weeks at the beginning and the end of the wait time target are the most congested, the policy resorts to overtime to free up some capacity for later arriving demand. It is worth noting that given the the way in which AOP policy books appointments it is highly unlikely that there will be any remaining demand that could be booked in those weeks. If there is insufficient capacity within the wait time target, the AOP policy first uses the overtime and then regular-time capacity, alternating between overtime and regular-time capacity on each week through the end of the horizon.

As regards follow-ups, the policy starts by booking appointments in the last week of the follow-up booking window and continues booking in reverse order from the last to the first week in the window, as the booking cost A_{ikn} gradually increases as we book down to the first week in the follow-up window. If there is insufficient capacity within the follow-up booking window, overtime is used (Figure 3.1b). The rule remain unchanged even if the follow-up wait time target is changed (Figure 3.2d) or additional follow-up visit (Figure 3.2e) is introduced to the model. If there is insufficient capacity within the follow-up booking window, the AOP policy alternates between overtime and regular-time capacity on each week after the window through end of the horizon.

Using the base-case scenario, we compare performance of the AOP policy against Myopic policy in a simulation with $R = 30$ replications. Each run has $T = 1500$ weeks of data collection period for performance metrics after $T_o = 5000$ weeks of warm-up period simulated using the DMB policy [33]. Numerical results of the simulation are summarized in Table 3.1. Results demonstrate that the AOP policy has cost savings potential through better service levels such as late bookings and average wait times resulting in significantly

lower average and discounted costs.

Table 3.1: Performance of the AOP Policy versus Myopic in Simulation

Criterion	**Urgency**	**AOP**	**Mypoic**
Average cost	-	32.42 ± 3.29	115.07 ± 13.91
Discounted cost	-	3368.25 ± 647.70	15751.80± 2632
Average capacity utilization	-	17.65 ± 0.03	17.86 ± 0.03
Percentage Late			
Consult	1	0.00 ± 0.00	0.50 ± 0.04
Consult	2	0.00 ± 0.00	0.09 ± 0.02
Follow-up	1	0.00 ± 0.00	0.00 ± 0.00
Follow-up	2	0.00 ± 0.00	0.00 ± 0.00
Average Wait Times			
Consult	1	1.70 ± 0.03	3.61 ± 0.17
Consult	2	5.27 ± 0.08	4.42 ± 0.22
Follow-up	1	12.00 ± 0.00	8.00 ± 0.00
Follow-up	2	12.00 ± 0.00	8.00 ± 0.00

It is important to note that the AOP policy shares some distinct and common characteristics with the PPQ policy for single visit setting [31] adapted to multiple visits setting we study here. We call the adapted version of PPQ policy a MDP heuristic 1 in the remainder of the chapter. First, the AOP policy books the consult appointments first in week 1, then in week $T_c(i)$ and working down to week 2 regardless of the priority class, whereas the MDP heuristic 1 books the high priority patients always as soon as possible. Second, the AOP policy uses overtime if there is no available regular-hour capacity in a given week within the wait time target unlike the MDP heuristic 1 which resorts to diversions/overtime only if there is no capacity within the entire wait time target. However, it is worth mentioning that the booking actions for the MDP heuristic 1 are constrained to regular-hour bookings and diversions only in week/day 1. The AOP and MDP heuristic 1 book follow-up appointments in a similar fashion and use overtime or diversions only if there is no available regular-hour capacity within the window.

We refer to the above set of booking rules derived from the AOP policy for small problem settings as the MDP heuristic 2 in the remaining of this chapter.

3.5.2 Case Study: An Endocrinology Outpatient Clinic

In order to evaluate the potential benefits that can be obtained from the use of the MDP heuristic 2, we will parameterize the simulation of the scheduling process using data from an

endocrinology clinic at a large teaching hospital in Ontario. The data include all referrals for consult for a cohort of patients who visited the clinic between Jan 1, 2015 and December 31, 2017. The clinic classifies patients into four priority classes with wait time targets of $14, 42, 90$ and 180 days for consult visits, respectively. The proportion of patients of each priority class is provided in Table 3.2.

Table 3.2: Priority Classes and Associated Wait Time Targets For Consults

Priority class	**Wait time target**	**Percent**
P1	14 Days	14%
P2	42 Days	25%
P3	90 Days	41%
P4	180 Days	20%

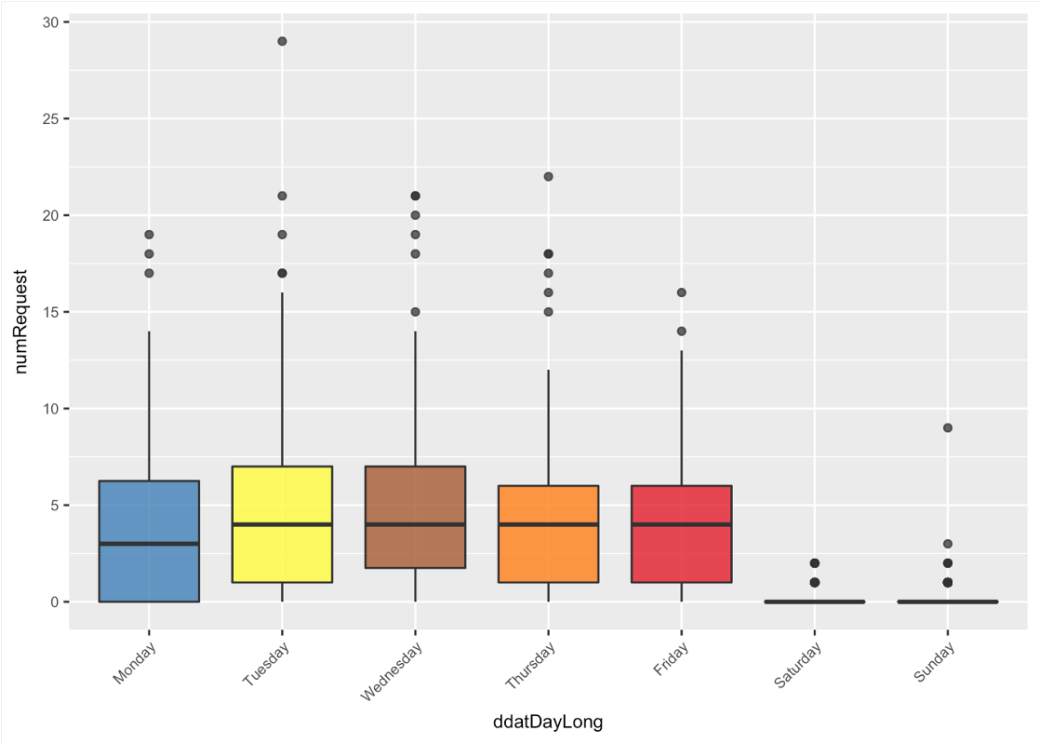

Figure 3.3: Distribution of Consult Requests by Day of the Week

A quick statistical analysis suggests that the demand for diabetes consults follows a Poisson distribution with the mean of 24 requests per week (p-value: < 0.0005). According to the data, most of the consult visits happened on business days while a small proportion happened on the weekends (Figure 3.3). Also, there does not seem to be a significant variation in demand among days of the week. While the diabetes clinic, on average, observed

5 requests per day on Tuesdays, Wednesdays and Thursdays, it received 4.5 requests per day on Mondays and Fridays.

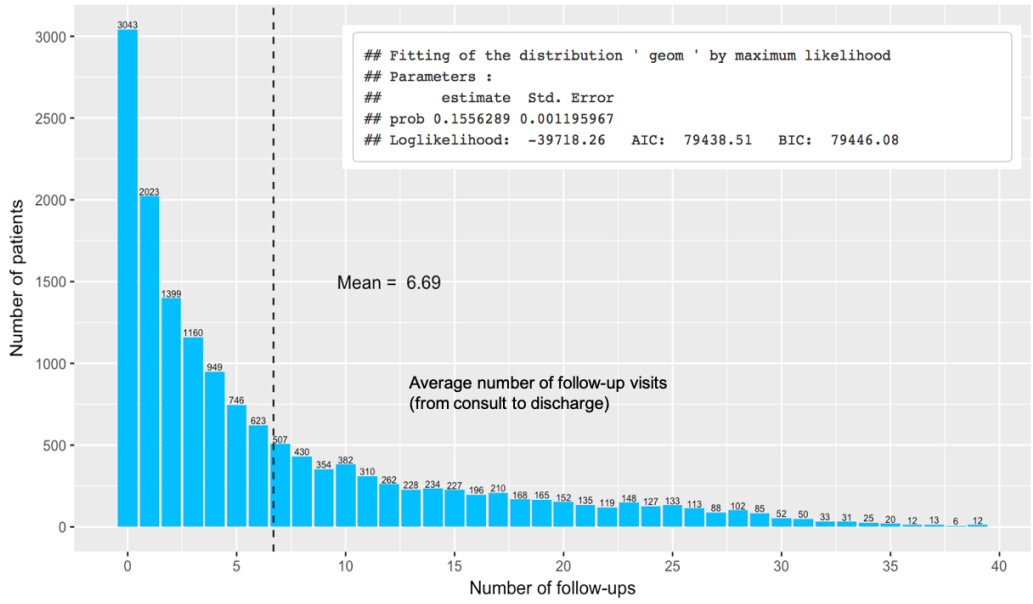

Figure 3.4: Distribution of the Number of Follow-up Appointments

Figure 3.4 shows the distribution of the number of follow-up visits performed after the consult. The data used includes all completed visits for a cohort of patients that visited the clinic between January 1, 2010 and December 31, 2016. If the time between two consecutive visits was more than 18 months then the appointment was coded as a consult. In addition, if a new physician was coded and the appointment showed as a consult then a new episode for the patient was created, otherwise the appointment was coded as follow-up. Only follow-up with physicians were included in our analysis. Dietitian or registered nurse visits were excluded.

The minimum number of follow-up visits after consult is zero, meaning that no follow-up visit was observed, whereas the maximum number of follow-up visits is 39. On average, 6.69 follow-up visits to diabetes clinics were required. We also note that the number of follow-up visits follows a Geometric distribution with parameter $p = 0.15$ meaning that probability of remaining in the system and requesting an additional visit is $p = 0.85$ or, alternatively, the probability of leaving the clinic after the current visit is $p = 0.15$.

Figure 3.5 shows the distribution of the number of days between visits. As it can be seen from the graph, most people are seen approximately every 180 days, with an average

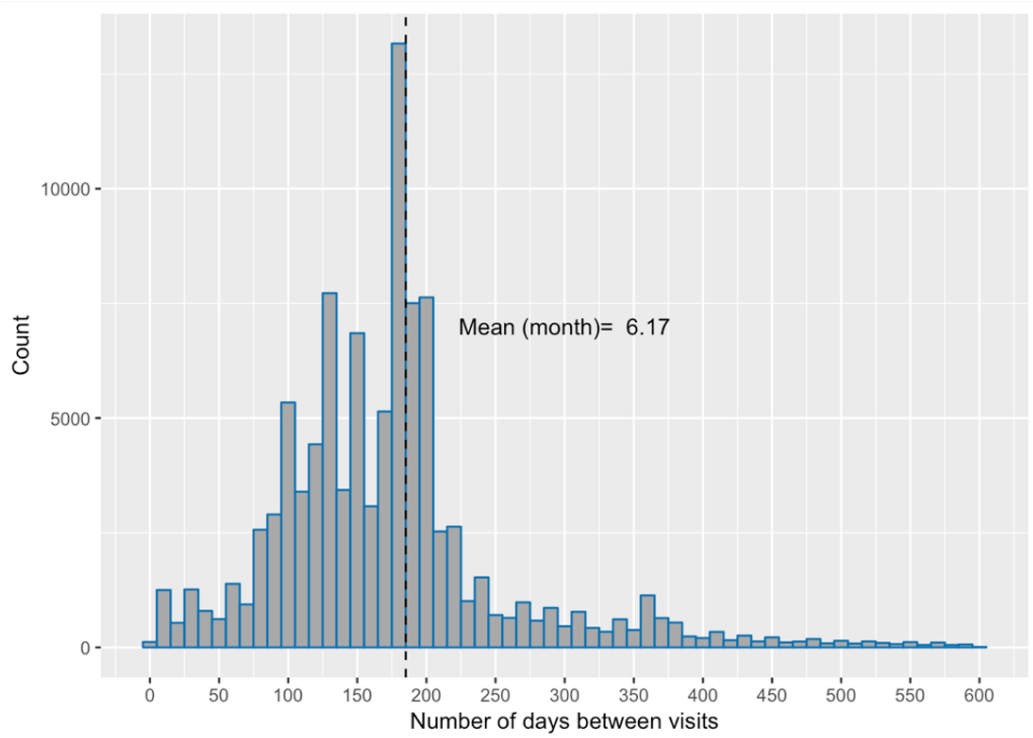

Figure 3.5: Number of Days Between Visits

of 6.17 months between visits.

3.5.3 Performance of the MDP Heuristic 2 in a Simulation

We compare the performance of the MDP heuristic 2 against Myopic, MDP heuristic 1 and DMB [33] policies in a simulation. We run each simulation with $R = 30$ replications and collect key performance statistics for $T = 1500$ weeks after $T_o = 5000$ weeks of warm-up period. We use the DMB booking policy for the warm-up period in all scenarios. The total time needed to simulate each scenario was in the range of 5 to 6 hours on a Windows 10 PC with Intel(R) Core (TM) i7 processor and 16.0 GB RAM. The simulation was implemented in Java version 13.0.1.

Table 3.3 summarizes the service level performance of the MDP heuristic 2 against three other policies. Service metrics for follow-up visits are averaged. In general, the MDP heuristic 2 outperforms the Myopic policy. However, it results in similar performance compared to the DMB and MDP heuristic 1 where the difference is not statistically significant. For priority 1 and 2 patient consult appointments, the MDP heuristic 2 provides 45% and 40% less average wait time compared to the Myopic policy. The Myopic policy books the priority

3 and 4 consult appointments much sooner than the MDP heuristic 2 since it chooses to book as soon as possible as opposed to considering the future impact of the booking decisions. The Myopic policy books 32%, 47% and 8% of priority 1, 2 and 3 consult patients visits, respectively, late, whereas the MDP heuristic 2, DMB and the MDP heuristic 1 do not book appointments late. For follow-up patient visits, mean wait time with the Myopic is around 24 weeks and with the other three policies it is 28 weeks. This observation is due to the fact that Myopic policy books follow-up visit as soon as possible while the MDP heuristic 2 and MDP heuristic 1 always book the follow-up visit on the last week within the window (26 ± 2). The DMB policy always seeks the week with the minimum number of bookings. In this setting, it is almost always going to be the last week that has the lowest utilization.

Table 3.3: Service Indicators with 95% Confidence Interval for Different Booking Policies

Perfomance Metric	Urgency	Visit Type	Myopic	MDP Heuristic 1	DMB	MDP Heuristic 2
Average Wait Times	1	consult	2.38 ± 0.04	1.16 ± 0.01	1.29 ± 0.06	1.31 ± 0.05
	1	follow-up	24.04 ± 0.00	28.00 ± 0.00	28.00 ± 0.00	28.00 ± 0.00
	2	consult	5.70 ± 0.05	3.49 ± 0.27	3.61 ± 0.44	3.42 ± 0.32
	2	follow-up	24.05 ± 0.00	28.00 ± 0.00	28.00 ± 0.00	28.00 ± 0.00
	3	consult	7.69 ± 0.07	9.82 ± 0.55	9.50 ± 0.77	9.27 ± 0.55
	3	follow-up	24.07 ± 0.00	28.00 ± 0.00	27.99 ± 0.00	28.00 ± 0.00
	4	consult	9.57 ± 0.09	24.34 ± 0.00	22.54 ± 0.55	23.86 ± 0.60
	4	follow-up	24.09 ± 0.00	28.00 ± 0.00	27.67 ± 0.00	28.00 ± 0.00
Proportion Late	1	consult	0.32 ± 0.01	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	1	follow-up	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	2	consult	0.47 ± 0.01	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	2	follow-up	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	3	consult	0.08 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	3	follow-up	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	4	consult	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	4	follow-up	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00

Figure 3.6 demonstrates the system-related key performance indicators. In general, the MDP heuristic 2, DMB and MDP heuristic 1 show similar performance. The total discounted cost is significantly higher with the Myopic policy versus the other three policies. The Myopic policy results in highest capacity utilization with 99.9% with the DMB policy being second highest at 99.7%. The Myopic resorts to approximately 4 overtime slots per week on average compared to the other policies which require less than 1.5 overtime slots. The total time to first available slot is equal among all four policies: 1 week for consult and 24 weeks for follow-up visits. The average cost per week for Myopic, MDP heuristic 1, MDP heuristic 2, and DMB policies are \$665.49 ± 50.21, \$127.22 ± 36.17, \$128.71 ± 31.86, and \$115.96 ± 32.27, respectively.

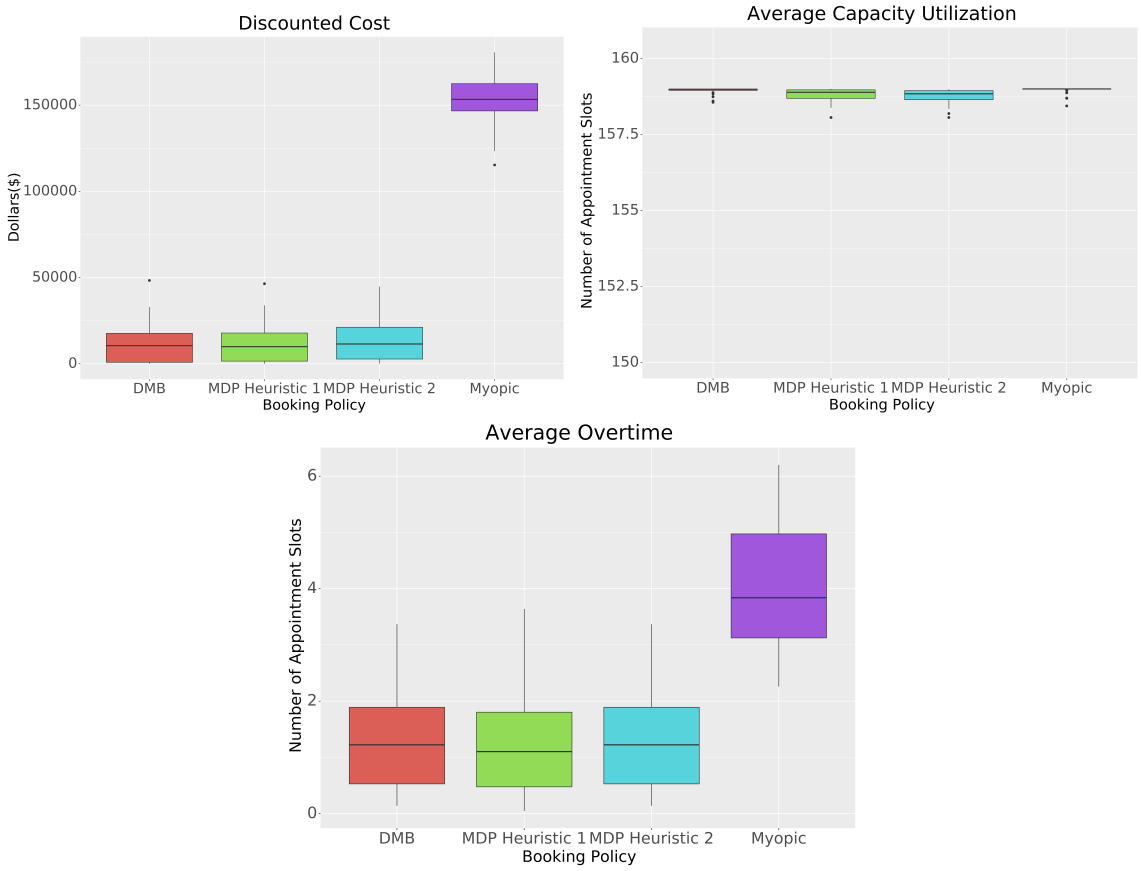

Figure 3.6: Performance of the Scheduling Policies in Terms of Different Metrics for the Endocrinology Clinic

3.6 Discussion

In this chapter, we tackle the follow-up appointment scheduling problem using approximate dynamic programming and derive a set of rules which can be used as booking guidelines in practice to reduce wait time times and achieve higher service levels with the least overtime capacity possible.

It is important to note that we present a generic dynamic advance scheduling model which is flexible enough to solve both the deterministic or stochastic number of follow-up visits. We illustrate the model utility using a large outpatient clinic setting where a follow-up visit occurs with a fixed probability estimated by a fitted Geometric distribution on real data. However, the model can handle any empirical distribution or probability mass function. Unlike some of the previously reported studies in the literature, we do not resort to a block scheduling approach and instead each appointment is booked individually better reflecting a real-life scenario.

We observe that the difference between the MDP heuristic 1 which is an extension of the PPQ policy [31] for a single visit to multiple visits setting and the MDP heuristic 2 is not practically or statistically significant. Thus, our model can be in a way seen as the extension of the PPQ policy. From a practical stand point, MDP heuristic 1 is somewhat easier to implement as there is no need to alternate between regular-time capacity and overtime is used only if there is insufficient capacity within the wait time target. In addition, the difference in performance between DMB and MDP heuristic 1 and 2 for multiple visit setting is not statistically or practically significant. This observation echoes the earlier finding for a single appointment scheduling setting [33].

In practice, booking clerks make the scheduling decisions based on their expertise and the urgency of the first or next visit for a particular patient. Considering that this task is repeated for many patients each day, it is a resource-intensive process with cognitive burden on clerks. The booking guidelines derived through the proposed approach enable booking clerks to make scheduling decisions that will ensure more efficient healthcare operations. However, manually applying them to scheduling operations on a daily basis can still be time- and resource-intensive. The process can be further streamlined if the booking guidelines are embedded in the engine of the appointment scheduling software used at a clinic integrated with the patient information system. Using the decision rules, the software can

automatically generate an appointment based on the existing bookings and new appointment requests with the associated urgency levels. This type of systems can improve the efficiency of the clinic operations.

For future direction, we plan to solve a special case of the model where the follow-up visits occur according to Geometric distribution. This is a simplified version of the model we solve here because there is no need to track the number of visits. In other words, the index k will be dropped from the state and action variables, and hence the entire MDP formulation. Our work can also be extended to include cancellations or no-shows and multiple types of resource capacities. In addition, different value function approximation architecture such as piece-wise linear value functions can be explored. We are also interested in developing closed-form analytical solutions for value function approximation parameters. Finally, we can extend the current model to determine the optimal panel size and number of different types of resources.

3.7 Bibliography

[1] Daniel Adelman and Adam J Mersereau. Relaxations of weakly coupled stochastic dynamic programs. *Operations Research*, 56(3):712–727, 2008.

[2] DW Alger. Appointment frequency versus treatment time. *American Journal of Orthodontics and Dentofacial Orthopedics*, 94(5):436–439, 1988.

[3] Ipsos Reid. Canadian Medical Association. Six in Ten Canadians Say Canadian Healthcare System is Not Sustainable Because of Costs, 2010.

[4] Davood Astaraky and Jonathan Patrick. A simulation based approximate dynamic programming approach to multi-class, multi-resource surgical scheduling. *European Journal of Operational Research*, 245(1):309–319, 2015.

[5] Dimitri P Bertsekas and John N Tsitsiklis. *Neuro-dynamic programming*. Athena Scientific, 1996.

[6] Donald G Boggs and M Anthony Schork. Determination of optimal time lapse for recall of patients in an incremental dental care program. *Journal of the American Dental Association (1939)*, 90(3):644–653, 1975.

[7] Tugba Cayirli and Emre Veral. Outpatient scheduling in health care: a review of literature. *Production and operations management*, 12(4):519–549, 2003.

[8] E Choan, Simone Dahrouge, Rajiv Samant, Ameneh Mirzaei, and Julie Price. Radical radiotherapy for cervix cancer: the effect of waiting time on outcome. *International Journal of Radiation Oncology* Biology* Physics*, 61(4):1071–1077, 2005.

[9] Daniela Pucci De Farias and Benjamin Van Roy. The linear programming approach to approximate dynamic programming. *Operations research*, 51(6):850–865, 2003.

[10] Francois d'Epenoux. A probabilistic production and inventory problem. *Management Science*, 10(1):98–108, 1963.

[11] Karen B DeSalvo, Bruce E Bowdish, Alys S Alper, David M Grossman, and William W Merrill. Physician practice variation in assignment of return interval. *Archives of internal medicine*, 160(2):205–208, 2000.

[12] Adam Diamant, Joseph Milner, and Fayez Quereshy. Dynamic patient scheduling for multi-appointment health care programs. *Production and Operations Management*, 27 (1):58–79, 2018.

[13] Yichuan Ding, Diwakar Gupta, and Xiaoxu Tang. Early reservation for follow-up appointments in a slotted-service queue. *Available at SSRN 3616565*, 2020.

[14] Viet Do, Val Gebski, and Michael B Barton. The effect of waiting for radiotherapy for grade iii/iv gliomas. *Radiotherapy and Oncology*, 57(2):131–136, 2000.

[15] G. Dobson, S. Hasija, and E. Pinker. Reserving capacity for urgent patients in primary care. *Production and Operations Researchs*, 20:456–473, 2011.

[16] A. Erdelyi and H. Topaloglu. Computing protection level policies for dynamic capacity allocation problems by using stochastic approximation methods. *IIE Transactions*, 41: 498–510, 2009.

[17] S Ayca Erdogan and Brian T Denton. Surgery planning and scheduling. *Wiley encyclopedia of operations research and management science*, 2010.

[18] Yasin Gocgun and Archis Ghate. Lagrangian relaxation and constraint generation for allocation and advanced scheduling. *Computers & Operations Research*, 39(10): 2323–2336, 2012.

[19] Yasin Gocgun and Martin L Puterman. Dynamic scheduling with due dates and time windows: an application to chemotherapy patient appointment booking. *Health care management science*, 17(1):60–76, 2014.

[20] L. Green and S. Savin. Reducing delays for medical appointments: A queuning approach. *Operations Research*, 56:1526–1538, 2008.

[21] L. Green, S. Savin, and B. Wang. Managing patient service in a diagnostic medical facility. *Operations Research*, 54:11–25, 2006.

[22] Diwakar Gupta and Brian Denton. Appointment scheduling in health care: Challenges and opportunities. *IIE transactions*, 40(9):800–819, 2008.

[23] Wei He, Yun Zheng, Ruhai Zou, Jingxian Shen, Junping Yang, Jiliang Qiu, Qiang Tao, Wenwu Liu, Zhiwen Yang, Yuanping Zhang, et al. Long-versus short-interval follow-up after resection of hepatocellular carcinoma: a retrospective cohort study. *Cancer Communications*, 38(1):26, 2018.

[24] Carlos Jackson, Mohammad Shahsahebi, Tiffany Wedlake, and C Annette DuBard. Timeliness of outpatient follow-up: an evidence-based approach for planning after hospital discharge. *The Annals of Family Medicine*, 13(2):115–122, 2015.

[25] Daniel L Kent, Robert A Nease, Harold C Sox JR, Linda D Shortliffe, and Ross Shachter. Eualuation of nonlinear optimization for scheduling of follow-up cystoscopies to detect recurrent bladder cancer. *Medical Decision Making*, 11(4):240–248, 1991.

[26] AG Leeftink, IA Bikker, IMH Vliegen, and RJ Boucherie. Multi-disciplinary planning in health care: a review. *Health Systems*, 9(2):95–118, 2020.

[27] Michael J Lichtenstein, Malcolm A Steele, Theodore P Hoehn, Christopher J Bulpitt, and Edward C Coles. Visit frequency for essential hypertension: observed associations. *Journal of Family Practice*, 28(6):667–673, 1989.

[28] MJ Lichtenstein, PM Sweetnam, and PC Elwood. Visit frequency for controlled essential hypertension: general practitioners' opinions. *The Journal of family practice*, 23 (4):331–336, 1986.

[29] N. Liu. Optimal choice for appointment scheduling window under patient no-show behaviour. 2013.

[30] Joren Marynissen and Erik Demeulemeester. Literature review on multi-appointment scheduling problems in hospitals. *European Journal of Operational Research*, 2018.

[31] Jonathan Patrick, Martin L Puterman, and Maurice Queyranne. Dynamic multipriority patient scheduling for a diagnostic resource. *Operations research*, 56(6):1507–1525, 2008.

[32] Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

[33] Antoine Sauré, Jonathan Patrick, and Martin L Puterman. Simulation-based approximate policy iteration with generalized logistic functions. *INFORMS Journal on Computing*, 27(3):579–595, 2015.

[34] Antoine Sauré, Jonathan Patrick, Scott Tyldesley, and Martin L Puterman. Dynamic multi-appointment patient scheduling for radiation therapy. *European Journal of Operational Research*, 223(2):573–584, 2012.

[35] Lisa M Schwartz, Steven Woloshin, John H Wasson, Roger A Renfrew, and H Gilbert Welch. Setting the revisit interval in primary care. *Journal of General Internal Medicine*, 14(4):230–235, 1999.

[36] Paul J Schweitzer and Abraham Seidmann. Generalized polynomial approximations in markovian decision processes. *Journal of mathematical analysis and applications*, 110 (2):568–582, 1985.

[37] Joanne K Tobacman, Rodney R Zeitler, Ana M Cilursu, and Motomi Mori. Variation in physician opinion about scheduling of return visits for common ambulatory care conditions. *Journal of general internal medicine*, 7(3):312–316, 1992.

[38] Van-Anh Truong. Optimal advance scheduling. *Management Science*, 61(7):1584–1597, 2015.

[39] H Gilbert Welch, Michael K Chapko, Kenneth E James, Lisa M Schwartz, Steven Woloshin, and Telephone Care Study Group. The role of patients and providers in the timing of follow-up visits. *Journal of general internal medicine*, 14(4):223–229, 1999.

[40] Siyun Yu, Vidyadhar G Kulkarni, and Vinayak Deshpande. Appointment scheduling for a health care facility with series patients. *Production and Operations Management*, 29(2):388–409, 2020.

4 Framework for Drug Formulary Decision Using Multiple-Criteria Decision Analysis[1]

In this chapter, we describe and apply a Multiple Criteria Decision Analysis (MCDA) approach to Canadian oncology drug reimbursement decisions. We use a recent extension of the UTilitiés Additives DIScriminantes approach, UTADIS^{GMS}. This approach deconstructs a portfolio of elements such as a set of drugs that have been reviewed and for which a decision has been made. It derives global and marginal utility functions that are consistent with the preferences exhibited by decision makers in their formulary decisions and determines the optimal utility threshold value which segregates the pre-defined ordered classes.

4.1 Introduction, Motivation and Related Literature

A formulary is a list of drugs that public or private insurance plans approve for use by their beneficiaries. A request to add a new drug to a formulary may come from different stakeholders such as manufacturers, patient groups, or clinicians [2]. The submitter forwards a file to a specialized review board or a pharmacy and therapeutics committee. The board or committee ultimately issues a reimbursement recommendation (i.e., full approval, conditional approval or rejection) based on the submitted file, on additional input from clinicians and patients, and on results from clinical and economic reviews.

We focus on the Canadian oncology drug review process. Provincial cancer agencies in Canada manage their oncology drug formularies based on recommendations from the pan-Canadian Oncology Drug Review (pCODR). New submissions undergo the pCODR review process and obtain a reimbursement recommendation before provincial plans and cancer agencies consider funding them. Detailed review reports are publicly available on the pCODR website [1].

[1]This chapter was published in Vusal Babashov, Sarah Ben Amor, and Gilles Reinhardt. Framework for drug formulary decision using multiple-criteria decision analysis. Medical Decision Making, 40(4):438447, 2020.

After accepting to review a submission, the board sets up an expert review team that evaluates the drug along four dimensions: clinical benefit, economic efficiency, patient-based values, and feasibility [2]. It then aggregates the information and formulates an initial recommendation, which becomes final subsequent to feedback from the stakeholders involved in the process.

This complex process takes months to complete and involves dozens of panellists, clinicians, scientists, and community members. They form a large pool of diverse expertise, and their input to the review board includes measurements, appraisals, and arguments derived from a multiplicity of sources. Aggregating and transforming that input into a recommendation is also challenging, since it involves debating over several independent and conflicting factors before reaching a consensus.

Systematic decision-support methods can facilitate the drug review process. For instance, several optimization-based approaches model listing recommendation decisions. Stinnett and Paltiel (1996) use a framework that maximizes effectiveness of a formulary subject to budgetary and ethical constraints [28]. Their model addresses some of the shortcomings of the standard cost-effectiveness analysis [35]. Olmstead and Zeckhauser (1999), inspired by the menu-setting approach, couch the formulary problem as a model that maximizes total consumer welfare or social surplus, subject to constraints on the total pharmaceutical budget over different patient groups and medical conditions [22]. Truong (2014) extends this approach to incorporate patient choice, stochastic utility and multiple drug categories and demonstrates that selecting a drug based on ICER generally results in suboptimal formularies [31]. However, concrete formulary listing decisions require thorough analysis on many criteria, which limits the applicability of optimization-based models.

Multiple Criteria Decision Analysis (MCDA) is a field of Operations Research that addresses decision-making problems that include conflicting objectives and criteria. It provides recommendations according to one of the following problematic: choice, ranking or sorting [24]. The sorting problematic concerns the assignment of the alternatives to pre-defined ordered classes, which is the goal of our study.

MCDA has gained traction in assessing health technologies and interventions [20]. A recent task force report from the International Society for Pharmacoeconomics and Outcomes Research (ISPOR) on emerging good practices in MCDA outlines different health-care settings in which MCDA tools can be applied, and provides guidelines on implementation

[29, 19]. As a result, HTA agencies around the world have started using MCDA [36, 9].

MCDA models traditionally use a direct approach in which criteria parameters, or *preferences*, are given. These models include the value system approach (in which *alternatives* are scored based on their criteria assessments), the outranking approach (based on pairwise comparisons of alternatives), and the multi-objective optimization approach (suitable when the set of alternatives is continuous rather than discrete) [30]. The disaggregation paradigm is an indirect approach that disentangles previous decisions. Unlike its direct counterparts, that paradigm infers a model that captures the preferences of the decision making bodies [25]. In so doing, it circumvents the challenges of scaling and weighting criteria that Health Technology Assessment (HTA) agencies have identified as a key concern with the implementation and interpretation of MCDA [21, 33].

A disaggregation approach is similar to a machine learning approach since both are data driven. Although they produce comparable results in many instances, there are important differences in terms of dimensionality, inconsistencies, and validation [8, 6]. Regarding dimensionality, machine learning operates with large data sets, whereas disaggregation approaches have sets with only a limited number of alternatives [34]. In our pCODR context, we note that committees consider a rather limited number of previous decisions, in addition to economic and clinical evidence, when issuing recommendations. We address the differences with respect to inconsistencies and validation in Sections 4.2 and 4.3.

Our goal in this chapter is to tailor a specific MCDA tool developed based on the disaggregation paradigm, the UTADIS $^{\mathrm{GMS}}$(UTilite´s Additives DIScriminantes, with the superscript referring to its authors)[11], and illustrate with a set of existing recommendations how it can positively contribute to the pCODR review process. UTADIS $^{\mathrm{GMS}}$ takes as input a set of submissions that have already been decided upon and criteria assessments, which may be qualitative or quantitative. It then infers parameters that model the contribution of each criterion considered in the analysis and produces a global utility value for the submissions that have been decided upon that is compatible with their respective recommendation.

Since recommendations are known in advance, the model analyzes how the criteria contribute to the recommendations and estimates a representative global utility function. Applying the function to the set of submissions yields the global utility value for each. The method also determines thresholds that partition the range of the utility values into

ordered classes that cluster recommendation categories for each alternative, for instance, "full approval," "conditional approval," and "rejection."

This method provides valuable decision making support. It ensures consistency of outcomes by implicitly inferring the contribution of each criterion to the recommendation and, in so doing, validates existing decisions. Moreover, we can leverage its predictive ability for alternatives without decisions with vectors of criteria assessments that model those potential new drugs.

We structure the remainder of this chapter as follows. In Section 4.2, we detail the UTADIS methodology and its extension, UTADIS$^{\text{GMS}}$. We also describe the data set we use in our analysis of the Canadian oncology drug reviews. In Section 4.3, we demonstrate the predictive capability of the model by developing 3 "scenarios" (potential new drugs) and analyzing the outcomes. We highlight our contributions, discuss the limitations of our model and findings, and point out avenues for future research in Section 4.4.

4.2 Methods

UTADIS was first presented by Devaud et al. (1980) as a criteria aggregation model to determine the classification of alternatives [7]. It is a version of the well established UTilites Additives method (UTA) used for ranking a set of alternatives from best to worst. An additive utility model is deemed sufficient to capture the decision maker's preferences when the criteria are independent [14]. UTADIS has been applied in financial decision making [18, 10], project and portfolio selection, [23] and business performance evaluation [32], notably in health care organizations [12]. Some extensions have also been developed [26, 37].

Greco et al. (2010) extended the original UTADIS method and built a multi-criteria sorting model based on U_{A^R}, the set of all global additive value functions [11]. Those are composed of monotonic and possibly nonlinear marginal value functions that are compatible with the provided preference information from the decision making bodies. This extension, UTADIS$^{\text{GMS}}$, is a classification method based on ordinal regression. It considers value functions that can be convex or concave rather than piece-wise linear. This so-called epsilon maximization approach, which we describe below, ensures that U_{A^R} is not empty. The solution to the linear program in this approach also provides the necessary and possible assignments of an alternative depending on whether its global utility is no less than that of

another alternative for all compatible value functions in U_{A^R} (necessary) or for at least one of them (possible).

We start with a set of alternatives with decisions, A^R, the *reference set*. We note that $A^R \subseteq A$ where $A = \{a_1, a_2, ..., a_n\}$ also includes alternatives without decisions [15]. Each alternative in A^R was previously evaluated over a set of m criteria, $G = \{g_1, g_2, \ldots g_j, \ldots, g_m\}$. We let X_j be the resulting vector of assessments on criterion j for the n alternatives and $x_j^1, x_j^2, \ldots x_j^{n_j(A)}$ be those assessments sorted in increasing order.

Drugs included in the reference set have already been classified into q ordered classes, $C_1, \ldots C_q$, where C_2 is preferred to C_1 and so on. In our context, we have $q = 3$ classes defined as "full approval" (C3), "conditional approval" (C2), and "rejection" (C1). Rather than using a direct procedure to obtain a global utility function for a given alternative, the preference disaggregation approach uses a regression-based technique to infer the function for that alternative. We assume the function to be additive and to exhibit the following structure:

$$U(a) = \sum_{j=1}^{m} u_j \left[g_j(a)\right] \in [0, 1] \tag{4.1}$$

where $U(a)$ is the global utility of alternative $a \in A$, and $u_j[g_j(a)]$ is the marginal utility of alternative $a \in A$ on evaluation criterion g_j. The global utility $U(a)$, scaled between 0 and 1, represents the evaluation of drug a over all m criteria. A drug will be classified into one of the predefined classes based on the resulting $U(a)$ and a set of utility thresholds $t = [t_1, t_2, \ldots, t_k, \ldots, t_{q-1}]$. Those thresholds, determined optimally, partition the q classes according to the following rule:

$$U(a) < t_1 \qquad a \in C_1 \tag{4.2}$$

$$t_k \leq U(a) < t_{k+1} \qquad a \in C_{k+1}, k \leq q-2 \tag{4.3}$$

$$t_{q-1} \leq U(a) \qquad a \in C_q \tag{4.4}$$

The solution to the following linear program yields the marginal utilities and thresholds.

$$\text{Max} \qquad \epsilon \tag{4.5}$$

subject to

$$U(a) \geq U(b) + \epsilon \qquad \forall a, b \in A^R : C(a) \geq C(b) \tag{4.6}$$

$$u_j(x_j^k) - u_j(x_j^{k-1}) \geq 0 \qquad j = 1, \ldots m, k = 2, \ldots, n \tag{4.7}$$

$$u_j(x_j^1) = 0 \qquad j = 1, \ldots m \tag{4.8}$$

$$\sum_{j=1}^{m} u_j(x_j^k) = 1 \qquad k = 1, \ldots n \tag{4.9}$$

$$U(a) \leq t_1 \qquad \forall a \in A^R : C(a) = 1 \tag{4.10}$$

$$t_k \leq U(a) \leq t_{k+1} \qquad \forall a \in A^R : C(a) = k, k = 1, \ldots, q-2 \tag{4.11}$$

$$t_{q-1} \leq U(a) \qquad \forall a \in A^R : C(a) = q \tag{4.12}$$

$$t_1 \leq \epsilon \tag{4.13}$$

$$t_k - t_{k-1} \geq \epsilon \qquad k = 2, \ldots q-1 \tag{4.14}$$

$$t_{q-1} \geq 1 - \epsilon \tag{4.15}$$

We note that for the objective function (4.5), a positive optimal value $\epsilon^* > 0$ ensures that U_{A^R} is not empty. Constraints (4.6) ensure that the resulting global value function is compatible with all pairs of decisions in set A^R. For instance, alternative a is assigned to a class that is no worse than the class to which alternative b is assigned if and only if the global utility of alternative a is no less than that of alternative b. Constraints (4.7) ensure the monotonicity of the resulting marginal utilities. Constraints (4.8) and (4.9) are normalization constraints that bound the marginal value functions between 0 and 1. Constraints (4.10) to (4.12) ensure that thresholds cluster alternatives in increasing order of classes. Given a positive optimal value $\epsilon^* > 0$, constraints (4.13) to (4.15) make the optimal thresholds distinct and within the [0,1] interval.

There are other solution approaches to determining value functions. Greco et al. (2011) introduced the concept of a representative value function from a set of compatible value functions to make robust sorting recommendations [11]. They obtain this function by maximizing the differences between the utilities of the alternatives across classes while minimizing the differences of the utilities within classes.

We implemented the UTADIS$^{\text{GMS}}$ method in R statistical software using the Rorutadis library [3], which is an implementation of the method with a set of value functions com-

patible with the decision making bodies' preference information. Skedgel et al. (2018) developed a data set of pCODR submissions reviewed between 2011 and 2017 [27]. With their permission, we use these data as input for the UTADIS ^{GMS} model to infer the global utility function and determine the classification thresholds. Subsequently, we conducted a scenario analysis by sampling from a set of compatible value functions to predict the assignment of submissions without recommendations (i.e., potential new drugs). Table 4.1 describes this process.

The data initially included 94 drug submissions with recommendations, 10 criteria, and the full, conditional, or rejection decision for each submission. Three records were removed from the 94 because of missing data. In the data consistency step, UTADIS^{GMS} removed 26 alternatives from the reference set to infer the utility model. An inconsistency occurs when, for a pair of alternatives, the decision of one conflicts with the other, given the criteria assessment values for each. As indicated in Table 4.1, the program identifies a minimal subset of alternatives that need to be removed so that the set of compatible pairs is free of conflicts and includes a maximum number of consistent alternatives [16]. The resulting reference set had 65 submissions with 11 full approval recommendations, 39 conditional approval recommendations, and 15 rejection recommendations. A drug may appear more than once if it was initially rejected and then resubmitted or if it was resubmitted for a different tumour type or indication.

The criteria are summarized in Table 4.2. We note that the relative survival gain is the ratio of the survival outcome measure (e.g., progression free, overall survival) of the treatment over that of a comparator. If there were no comparators, the ratio was set to 1. The quality of the clinical evidence criterion is a binary variable indicating that the submission was based on a two-arm double-blind randomized phase III clinical trial. The severity of adverse reactions (compared to existing treatments) is also binary, with 1 indicating lower adverse reactions. The incremental cost-effectiveness ratio(ICER) is a continuous scale variable. The quality of the ICER can be high (= 1) or low (= 0) and indicates confidence in the ICER estimate. The unmet needs criterion is equal to 1 if there are no alternatives. The budget impact is another categorical variable that can take a high (= 1) or low (= 0) value. To make our model parsimonious and our results statistically robust, we excluded the overall survival flag, infrastructure, and type of drug attributes, since those were correlated with the seven retained criteria or were not sufficiently

Table 4.1: Description of the UTADIS^{GMS} Modelling Process

Phase	Step	Output	Notes
Structuring the problem	Identification of criteria [27]	Of the 10 criteria identified, 7 are retained in our analysis.	We kept criteria that are most relevant (i.e., independent and discriminating in terms of preferences). See Table 4.2
	Tabulation of alternatives	There were 91 alternatives in their final sample size.	
	Measurement of Performance [27]	The performance matrix includes criteria assessments and decision for each of the 91 alternatives.	See [5].
	Scenario Development	We developed 3 scenarios (alternatives without a decision) and appended their criteria assessments to the set of alternatives.	See Table 4.2.
Applying UTADIS^{GMS}	Data consistency	Inconsistencies are identified.	Pairs of alternatives with conflicts are identified by brute force technique & *rorutadis*.
		Alternatives causing inconsistencies are removed.	See [5].
	Generation of results	The output includes: 1. The marginal utility functions for each criterion. 2. The thresholds that partition the three classes. 3. The global utility value and class assignment for all alternatives, including the scenarios.	See Figure 4.1. Also, see [5].
Analysis and model validation	Cross-validation	Thresholds and classification for each testing sample.	See [5].
	Robustness of scenario results	Percentage of value functions for which scenarios remain in the same class assignment.	Stochastic results are obtained with the *CalculateStochasticResults* function from *rorutadis*.

discriminating in terms of preferences.

Table 4.2: Evaluation Criteria and Scenario Values

Evaluation criteria	Definition	Scenario 1	Scenario 2	Scenario 3
Relative survival gain	Survival gain versus comparator (continuous, range = [0.8, 5.3])	1.1	1.5	1
Quality of clinical evidence	Quality of clinical study with respect to phase of the random clinical trial and results (binary, 1 = high quality, 0 = low quality)	1	1	0
Severity of adverse reactions	Severity of adverse reactions compared to existing treatment (binary, 1 = lower adverse reactions, 0 = higher adverse reactions)	1	0	1
ICER	Incremental Cost Effectiveness Ratio ($K, continuous, range = [18, 461])	75	150	200
ICER Quality	Uncertainty of ICER, Sensitivity analysis (binary, 1 = high uncertainty, 0 = low uncertainty)	0	0	1
Unmet needs	(binary: 1 = existence of alternatives, 0 = absence of alternatives)	1	1	0
Budget impact	Estimated impact based on the patient population size and available treatments (binary, 1 = high impact, 0 = low impact)	1	0	0
Overall survival flag	Not included			
Infrastructure	Not included			
Type of drug	Not included			

4.3 Results

Figure 4.1 reports the output of the UTADIS$^{\mathrm{GMS}}$ model with the global utility values for each drug submission, sorted in increasing order. The two dashed vertical lines indicate the optimal threshold values between the rejected and conditional approval classes (0.316) and between the conditional and full approval classes (0.632).

The last three bars in Figure 1 report the global utility values and the classes of three scenarios (potential new drugs) which vectors of criteria assessments are provided in Table 2. We label the hypothetical drugs "incremental" (Scenario 1), "drastic" (Scenario 2), and "orphan" (Scenario 3). Scenario 1 models a potential incremental drug, with a relative survival gain of 1.1 (or 10% over an existing comparator). We further assume that it has a high quality clinical evidence, that it causes lower adverse reactions, and that alternative options already exist. We set its ICER at $75,000, which corresponds to the first quartile (i.e., 25th percentile) values for the 65 submissions, and a high-budget impact that would result from high demand. Rorutadis estimates the global utility value of this potential alternative at 0.421. Since it lies between the two thresholds $(0.316, 0.632)$, it falls into the conditional approval class.

Scenario 2 models a potential drastic drug that is more efficacious with a relative survival gain of 1.5. We assume it exhibits a higher ICER than the incremental drug ($150,000) and higher adverse reactions compared to its existing alternatives. Contrary to the incremental drug in Scenario 1, we set the budget impact at "low" based on lower expected demand. Its

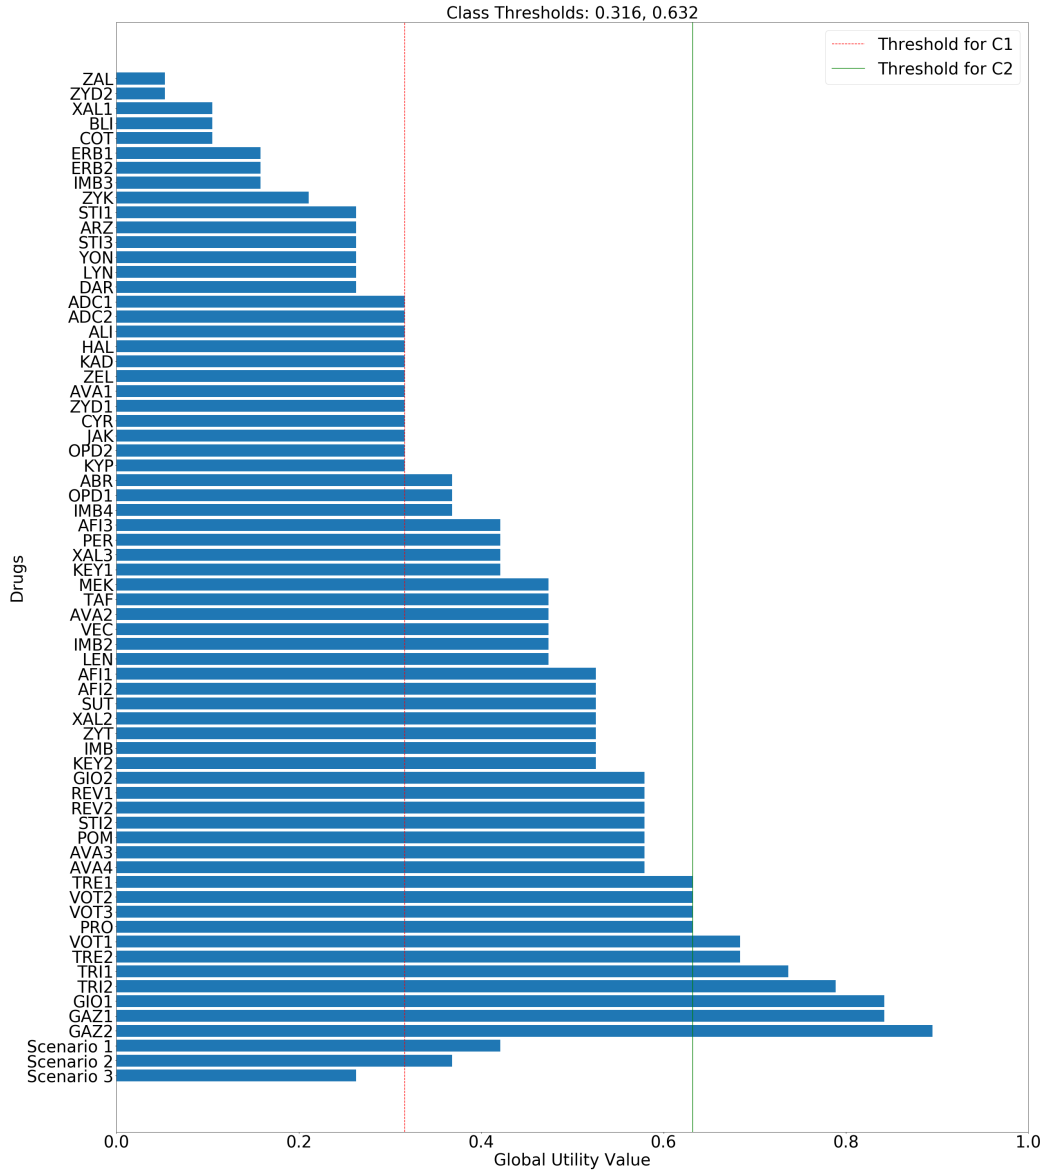

Figure 4.1: Global (Comprehensive) Utility Values and Classification of Alternatives

global utility returned at 0.368 and it is thus also categorized in the conditional approval class.

Scenario 3 models a potential orphan drug. We set its relative survival gain at 1.0, since there is no comparator, and the quality of clinical evidence at "low." We assume it is at phase II of a single-arm clinical trial and that it will cover unmet therapeutic needs. We further assume low adverse reactions, no alternative treatment, and a low budget impact due to low demand. We set its ICER at $200,000, which corresponds to the third quartile (i.e, 75th percentile) values of the reference set. We also assume its ICER exhibits high

uncertainty, which is typical of orphan drugs. The resulting global utility value is 0.263, which puts it in the rejected class.

In a disaggregation approach, fit statistics and related error measures are moderately useful for validation purposes. What is central to validation is the interaction with the decision makers [8, 6]. Since it is their decision process that is being represented, validation consists of informing them of the model output. They can then analyze this information to review or adjust the decision criteria and reference set. We nevertheless performed cross-validation by applying the model 5 times with distinct testing samples of $n = 5$ alternatives pulled from the reference set and another 5 times with testing samples of $n = 10$ alternatives. In the $n = 5$ tests, we noted that only 1 alternative in one sample was categorized differently. In the $n = 10$ tests, we noted 2 samples with 4 alternatives categorized differently, 2 samples with 2 alternatives categorized differently, and 1 with a single alternative categorized differently. Such results illustrate the impact of eroding information from the increasing test samples on the model output. Less data implies a higher propensity of bias and a diminished learning capacity of the model. This further justifies interaction with the committee to determine the appropriate size of the reference set for this context.

We performed a validation exercise more in line with a disaggregation approach by applying the *CalculateStochasticResults* Rorutadis function. It tested the 3 scenarios against a randomly generated sample of 1000 representative functions. It assigned the incremental drug to the conditional approval class for 89% of the representative functions and to the full approval class for 11%. It assigned the drastic drug to the conditional approval class for 100% of the functions and the orphan drug to the rejected class for 100% of the functions. These percentages show that the model is robust since the assignments are almost the same for the sample of 1000 representative functions.

4.4 Discussion

In this article, we apply a methodological framework from MCDA to investigate the potentials and applicability of a preference disaggregation approach to drug formulary design. The main advantage of this approach is that it infers preferences from the data as opposed to other MCDA methods that rely on direct assessments of model parameters and con-

sequently on a greater cognitive effort from the decision making bodies [15]. In addition to objectively sorting and determining thresholds for the inclusion of existing drugs, this method can be applied to a new drug using estimates or targets for the criteria assessments to help predict its recommendation outcome. This is useful information that can assist review boards in their decision process. It can also help research and development divisions of companies to prioritize resources and innovation capabilities.

The method have been successful in numerous fields, so most of the limitations of this work are related to the formulary application. We included criteria and alternatives based on the Skedgel et al. (2018) analysis [27]. We could add criteria that explicitly assess clinical judgment, demand, patient-based values, or adoption feasibility. We could also expand the binary criteria scales to reflect measurements that are more accurate. For instance, the range for the quality of the clinical evidence criterion could be widened in light of more precise information. In addition, although the committee reviews dozens of submissions each year, the model does not consider the timing of the recommendations. Thus, as new submissions and decisions become available, we could update the reference set to reflect a more current or precise set of alternatives that captures the evolving preferences of the committee. For instance, we could truncate older submissions from the set or partition it into subsets defined by tumour type or indication.

In the MCDA literature, disaggregation methods are all based on additive utility functions. When the criteria are independent, this is sufficient to capture the decision makers' preferences [14]. Independence is a reasonable assumption in the pCODR context, as shown in Table 4.2. In more complex decision making settings, an additive utility function may not properly capture interactions between criteria. A more general form could possibly be developed and tested to model preferences.

Our model was proved to be robust based on the results obtained from applying the *CalculateStochasticResults* Rorutadis function. In MCDA, robustness refers to the analysis of imperfections of the model. Assessing robustness consists of verifying whether the provided conclusions match those from a randomly generated set of compatible utility functions. Experimental work shows that the robustness of an additive model depends on both the number of criteria and the size of the reference set [17].

The approach allows for the identification of inconsistent decisions. This resulted in the removal of 26 alternatives from the reference set to infer the utility model. Inconsistencies

can be explained by many factors. In decision making, criteria not explicitly stated can nonetheless influence outcomes. Also, perceptions and preferences of the decision making body may change over time and context. For example, the assessment of the budget impact may depend on the economic context in which the review occurs. Moreover, committee membership turnover may influence the way it processes information and makes recommendations. This limits the overall "memory" of the committee and does not ensure that the recommendation of a new submission is consistent with previous ones. In addition, committee members, based on their own field of expertise, may favour submissions that affect some patient groups, conditions, or tumour types over others [27]. To control inconsistencies, the disaggregation approach calls for continuous interaction and feedback with the committee [6]. This exercise helps calibrate the model and ultimately lowers the number of inconsistent decisions.

Disaggregation analysis does not replace but rather facilitates the decision support process. It helps decision makers gain insight into the problem's alternatives and criteria, the modelling process, and the interpretation of output. The inferred model provides a starting basis for the decision making support process. If the model output agrees with the preferential system of the decision makers, then the model can be directly applied to new decision instances. Otherwise, the decision makers can revise their criteria and reference set or provide recommendations about the calibration of the model. This approach provides an efficient, quick, and reliable way to support the review process that ensures inclusion of relevant criteria and consistency among decisions.

In a wider analysis, more therapeutic conditions can be modelled in the set of alternatives of the formulary. Moreover, formularies can be multi-tiered, with each tier defined by a cost-sharing policy, which is typical of private insurance plans. Insurers and pharmacy benefit managers also design formularies for a multiplicity of plans. Thus, formulary candidates can be modelled as alternatives by developing appropriate sets of criteria. The methodology above could then be extended and applied to glean strategic value of candidate formularies, not only to insurers but also to employers or other clients who, in turn, provide coverage options to their workforce or constituents.

This article contributes to an ongoing discussion on the use of MCDA for HTA decisions, [15, 4, 13] specifically the critical value of proven and tested methodologies for evaluating treatment options. We focused on recent extensions to the UTADIS method and expanded it

to model the reimbursement recommendation decisions of new or existing prescription drugs. We applied it using data on Canadian oncology drugs reviewed between 2011 and 2017. By associating the input and requirements of the method with the information germane to the offerings, we obtained valuable information on their potential outcome in terms of global utility value and class. We also developed scenarios that illustrate how the method can be applied for predictive purposes.

4.5 Bibliography

[1] Find a Review.cadth pan-Canadian Oncology Review. Accessed: October 15, 2018. URL `https://www.cadth.ca/pcodr/find-a-review`.

[2] Process in Brief. Canadian Agency for Drugs and Technologies in Health. Accessed: May 31, 2019. URL `https://www.cadth.ca/about-cadth/what-we-do/products-services/cdr/common-drug-reviewsubmissions/process-in-brief`.

[3] Robust Ordinal Regression Utadis. january 17, 2017. URL `https://cran.r-project.org/web/packages/rorutadis/rorutadis.pdf`.

[4] Aris Angelis and Panos Kanavos. Multiple criteria decision analysis (mcda) for evaluating new medicines in health technology assessment and beyond: the advance value framework. *Social Science & Medicine*, 188:137–156, 2017.

[5] Vusal Babashov, Sarah Ben Amor, and Gilles Reinhardt. Framework for drug formulary decision using multiple-criteria decision analysis. *Medical Decision Making*, 40(4):438–447, 2020.

[6] Salvatore Corrente, Salvatore Greco, Miłosz Kadziński, and Roman Słowiński. Robust ordinal regression in preference learning and ranking. *Machine Learning*, 93(2-3):381–422, 2013.

[7] JM Devaud, G Groussaud, and Eric Jacquet-Lagreze. Utadis: Une méthode de construction de fonctions d'utilité additives rendant compte de jugements globaux. *European Working Group on Multicriteria Decision Aid, Bochum*, page 94, 1980.

[8] Michael Doumpos and Constantin Zopounidis. Preference disaggregation and statistical learning for multicriteria decision support: A review. *European Journal of Operational Research*, 209(3):203–214, 2011.

[9] Dóra Endrei, Bálint Molics, and István Ágoston. Multicriteria decision analysis in the reimbursement of new medical technologies: real-world experiences from hungary. *Value in Health*, 17(4):487–489, 2014.

[10] Chrysovalantis Gaganis, Fotios Pasiouras, and Constantin Zopounidis. A multicriteria

decision framework for measuring banks' soundness around the world. *Journal of Multi-Criteria Decision Analysis*, 14(1-3):103–111, 2006.

[11] Salvatore Greco, Vincent Mousseau, and Roman Słowiński. Multiple criteria sorting with a set of additive value functions. *European Journal of Operational Research*, 207 (3):1455–1470, 2010.

[12] Evangelos Grigoroudis, Eva Orfanoudaki, and Constantin Zopounidis. Strategic performance measurement in a healthcare organisation: A multiple criteria approach based on balanced scorecard. *Omega*, 40(1):104–119, 2012.

[13] William Hall, Craig Mitton, and Adrian Levy. Realizing the "great potential" of mcda in hta. *Value in Health*, 21(12):1355–1356, 2018.

[14] Eric Jacquet-Lagreze and Jean Siskos. Assessing a set of additive utility functions for multicriteria decision-making, the uta method. *European journal of operational research*, 10(2):151–164, 1982.

[15] Eric Jacquet-Lagreze and Yannis Siskos. Preference disaggregation: 20 years of mcda experience. *European Journal of Operational Research*, 130(2):233–245, 2001.

[16] Miłosz Kadziński, Marco Cinelli, Krzysztof Ciomek, Stuart R Coles, Mallikarjuna N Nadagouda, Rajender S Varma, and Kerry Kirwan. Co-constructive development of a green chemistry-based model for the assessment of nanoparticles synthesis. *European journal of operational research*, 264(2):472–490, 2018.

[17] Milosz Kadzinski, Mohammad Ghaderi, Jakub Wkasikowski, and Nuria Agell. Expressiveness and robustness measures for the evaluation of an additive value function in multiple criteria preference disaggregation methods: An experimental analysis. *Computers & Operations Research*, 87:146–164, 2017.

[18] Kyriaki Kosmidou and Constantin Zopounidis. Predicting us commercial bank failures via a multicriteria approach. *International Journal of Risk Assessment and Management*, 9(1-2):26–43, 2008.

[19] Kevin Marsh, Maarten IJzerman, Praveen Thokala, Rob Baltussen, Meindert Boysen, Zoltán Kaló, Thomas Lönngren, Filip Mussen, Stuart Peacock, John Watkins, et al.

Multiple criteria decision analysis for health care decision making—emerging good practices: report 2 of the ispor mcda emerging good practices task force. *Value in health*, 19(2):125–137, 2016.

[20] Kevin Marsh, Tereza Lanitis, David Neasham, Panagiotis Orfanos, and Jaime Caro. Assessing the value of healthcare interventions using multi-criteria decision analysis: a review of the literature. *Pharmacoeconomics*, 32(4):345–365, 2014.

[21] Kevin D Marsh, Mark Sculpher, J Jaime Caro, and Tommi Tervonen. The use of mcda in hta: great potential, but more effort needed. *Value in Health*, 21(4):394–397, 2018.

[22] Todd Olmstead and Richard Zeckhauser. The menu-setting problem and subsidized prices: drug formulary illustration. *Journal of Health Economics*, 18(5):523–550, 1999.

[23] Konstantina Pendaraki, Constantin Zopounidis, and Michael Doumpos. On the construction of mutual fund portfolios: A multicriteria methodology and an application to the greek market of equity mutual funds. *European Journal of Operational Research*, 163(2):462–481, 2005.

[24] Bernard Roy. *Multicriteria methodology for decision aiding*, volume 12. Springer Science & Business Media, 2013.

[25] Yannis Siskos and Athanasios Spyridakos. Intelligent multicriteria decision support: Overview and perspectives. *European Journal of Operational Research*, 113(2):236–246, 1999.

[26] Yannis Siskos and Denis Yannacopoulos. Utastar: An ordinal regression method for building additive value functions. *Investigaçao Operacional*, 5(1):39–53, 1985.

[27] Chris Skedgel, Dominika Wranik, and Min Hu. The relative importance of clinical, economic, patient values and feasibility criteria in cancer drug reimbursement in canada: a revealed preferences analysis of recommendations of the pan-canadian oncology drug review 2011–2017. *Pharmacoeconomics*, 36(4):467–475, 2018.

[28] Aaron A Stinnett and A David Paltiel. Mathematical programming for the efficient allocation of health care resources. *Journal of Health Economics*, 15(5):641–653, 1996.

[29] Praveen Thokala, Nancy Devlin, Kevin Marsh, Rob Baltussen, Meindert Boysen, Zoltan Kalo, Thomas Longrenn, Filip Mussen, Stuart Peacock, John Watkins, et al. Multiple criteria decision analysis for health care decision making—an introduction: report 1 of the ispor mcda emerging good practices task force. *Value in health*, 19(1): 1–13, 2016.

[30] Praveen Thokala and Alejandra Duenas. Multiple criteria decision analysis for health technology assessment. *Value in Health*, 15(8):1172–1181, 2012.

[31] Van-Anh Truong. Optimal selection of medical formularies. *Journal of Revenue and Pricing Management*, 13(2):113–132, 2014.

[32] Fotini Voulgaris, Michael Doumpos, and Constantin Zopounidis. On the evaluation of greek industrial sme's performance via multicriteria analysis of financial ratios. *Small business economics*, 15(2):127–136, 2000.

[33] Andrew Walker. Challenges in using mcda for reimbursement decisions on new medicines? *Value in Health*, 19(2):123–124, 2016.

[34] Jyrki Wallenius, James S Dyer, Peter C Fishburn, Ralph E Steuer, Stanley Zionts, and Kalyanmoy Deb. Multiple criteria decision making, multiattribute utility theory: Recent accomplishments and what lies ahead. *Management science*, 54(7):1336–1349, 2008.

[35] William S Weintraub and David J Cohen. The limits of cost-effectiveness analysis. *Circulation: Cardiovascular Quality and Outcomes*, 2(1):55–58, 2009.

[36] Sitaporn Youngkong, Rob Baltussen, Sripen Tantivess, Adun Mohara, and Yot Teerawattananon. Multicriteria decision analysis for including health interventions in the universal health coverage benefit package in thailand. *Value in health*, 15(6):961–970, 2012.

[37] Constantin Zopounidis and Michael Doumpos. Prefdis: a multicriteria decision support system for sorting decision problems. *Computers & Operations Research*, 27(7-8):779–797, 2000.

5 Concluding Remarks

In this thesis, we take an in-depth look at three practical healthcare problems. In Chapter 2, we study wait time targets and present a framework to set wait time targets for multi priority patient settings. We start with a discussion about a practical way of scheduling patients in a cost-efficient manner. Then, we simulate the chosen policy and estimate key service performance measures such as the average wait time and overtime utilization. Next, we regress these two metrics on wait time targets to estimate the corresponding functions. Finally, we use inverse optimization to determine the unit cost of waiting for a chosen set of wait time targets and the associated total cost to the system. We illustrate the benefits of the proposed approach by applying it to a case study based on data from a rhemuatology clinic.

In Chapter 3, we solve a dynamic capacity allocation problem that involves multiple priority patients with the potential requirement for one or more follow-up appointments. We develop a generic model that is able to handle both a fixed and a probabilistic number of visits where the probability for a specific visit can follow a theoretical distribution (e.g., Geometric or an empirical distribution). Using data from an endocrinology clinic, we demonstrate significant cost savings that could potentially be achieved if the proposed heuristic is implemented in this type of settings.

In Chapter 4, we study a drug formulary design problem with a preference-disaggregation-based paradigm method. We illustrate this approach using an open-source decision support tool and published data which can be useful for policy makers to make quick and efficient budget allocation decisions for pharmaceutical drug formulary decisions.

All three chapters of the thesis make a modelling contribution to the healthcare operations research analytics community. Also, the models and numerical results have the potential to impact health practice by informing policy and patient care related decisions. We use a range of analytical tools to solve these three practical problems using prescriptive methods such as linear program, mixed integer program, simulation, Markov decision pro-

cess, Approximate Dynamic Programming, multi-criteria decision analysis and predictive tools such as regression and deep neural networks.

The efficient use of medical resources is crucial for achieving timely, high quality and cost-effective care. By developing, implementing and utilizing advanced analytics methodologies, the first two papers in the thesis provide systematic ways of identifying effective resource allocation guidelines for patients with different urgency levels. These guidelines could be used in practice to significantly reduce wait times for health care services and thus to potentially decrease the impact of delays on patients' health. This research has also the potential to provide valuable insights into the long-term workload incurred by new patient admissions and the impact of follow-up visits on patient wait times and physician utilization. Thus, it can help administrators make crucial decisions that will ensure timely access to health care for patients and more cost-effective care delivery for health care organizations. The last paper provides a more robust approach to HTA decisions to ensure a cost-effective allocation of budget and a more holistic and coherent approach to reimbursement decisions.