Dynamic Robust Multi-Class Advance Patient Scheduling

by

Hamidreza Khajeh Arzani

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Management Sciences

Waterloo, Ontario, Canada, 2022

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this work, we study an advance patient scheduling problem where patients of different classes have different service times and incur different waiting costs to the system. It is known in the literature that multi-class advance dynamic patient scheduling is a challenging problem due to the high variability in the daily arrival process of patients, as well as the high dimensionality of the problem. To overcome these challenges, we develop a novel dynamic optimization framework where the multi-class advance scheduling problem can be approximately decomposed to multiple single-stage stochastic programs. Furthermore, we develop a distributionally robust formulation and quantify uncertainty in arrivals by applying a risk-averse optimization approach. Exploiting patient-level offline data, we develop a data-driven algorithm to minimize the worst-case outcome that may happen due to the high variability in arrivals. We examine the performance of the proposed robust algorithm by leveraging the MRI data from hospitals in Ontario and show that the dynamic robust model outperforms the dynamic stochastic approach significantly. We also observe that the proposed robust model performs well compared to an offline policy, which is based on the full knowledge of the future arrivals.

Acknowledgements

My sincerest thanks to my advisor, Prof. Hossein Abouee-Mehrizi, for the endless guidance, support, understanding, patience, and encouragement. He is a very good exemplar of a researcher, a mentor, and an instructor, and always inspires me to think more logically and precisely. Professor Mehrizi always took the extra step to show his care to my current and future success. Working with him has been an invaluable experience for me.

A special thanks to Prof. Saeed Ghadimi for being my unofficial advisor and giving me valuable comments and suggestions. I could not have undertaken this journey without him, who generously provided knowledge and expertise. I also would like to thank the reader of this thesis, Prof. Mehrdad Pirnia, for his valuable comments and guidance along the way.

A warm thank you to all my friends in Waterloo who always supported me and helped me in many different ways. Thank you all for making my time at University of Waterloo so memorable. Finally, I would like to express my deep appreciation to my parents, Rasoul and Mina, and my dear wife, Ida, for their endless support, patience, understanding and their true love.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Reduced capacity for elective medical procedures during the Covid-19 pandemic has created significant backlogs. The CovidSurg Collaborative (CSC) estimated in June 2020 that, due to COVID-19, more than 28 million elective surgeries would be canceled or delayed during the 12 weeks of disruption worldwide. In England, more than 5.1 million operations are waiting to be performed [1]. Also, the province of Ontario in Canada faces a backlog of over 15 million procedures with MRI services having the highest (around 500,000) number of canceled or delayed procedures [2]. The goal of the healthcare policymakers is to reduce the backlogs and bring down the wait times to acceptable levels. For example, the government of Ontario recently decided to allocate 324 million dollars[3] and run hospitals at 115% capacity to clear pandemic surgical backlogs[4]. A key question that health policymakers and clinical administrators are facing nowadays is that what strategy they should adopt for scheduling patients without undue significant additional costs.

In this paper, motivated by Ontario MRI services, we study a multi-class advance patient scheduling problem where patients from different classes arrive sequentially over a finite time horizon. The classes of patients are distinguished based on patients' priorities, service times, and arrival distributions.

There are two main challenges in dealing with multi-class advance scheduling problems.

[1]https://www.economist.com/britain/2021/07/10/the- mystery-at-the-heart-of-the-national-health-service

[2]https://www.cbc.ca/news/health/oma-covid-pandemic-medical-procedure-backlog-1.6059285

[3]https://www.theglobeandmail.com/canada/article-ontario-to-spend-324-million-to-handle-surgery-backlog-left-by-covid/

[4]https://www.cbc.ca/news/canada/toronto/covid-19-ontario-july-28-2021-surgical-backlog-recovery-1.6120588

First, using the standard dynamic programming to derive the optimal scheduling policy is not straightforward due to the high dimensionality. In fact, it is known in the literature that the advance scheduling problem suffers from the curse of dimensionality (see e.g., [47], [60], [11]). Second, in practice, the variability in patients' daily arrivals can be extremely high as indicated in [38], [47] and [60]. Therefore, an efficient scheduling policy should lead to *robust* scheduling decisions and protect the system against the high variability in the daily arrivals. Our goal in this work is to address these two challenges by developing a novel dynamic optimization framework and proposing a robust online data-driven algorithm to solve the advance patient scheduling problem.

To overcome the curse of dimensionality, we develop a novel dynamic optimization framework. Specifically, we modify the multi-class advance scheduling problem in two major steps. First, we formulate the problem as a multi-stage stochastic program and decompose it to multiple single-stage stochastic programs by relaxing the *non-anticipativity* constraints. Intuitively, on each day, instead of solving a dynamic stochastic problem, we consider a static model and schedule patients who have arrived while reserving capacity for patients who will arrive in the following days. Reserving capacity for future arrivals would reduce the waiting time of the high priority patients that should be served relatively fast. In the second step, we reformulate the feasible set associated with the scheduling of future arrivals. This reformulation adds more flexibility to each single-stage stochastic program to control the uncertainty in arrivals. Finally, we propose a data-driven algorithm to solve the single-stage stochastic program on each day. The outline of the sub-problem that should be solved day-to-day is presented in Figure 1.1. Generally speaking, on each day, the demand is first realized and then, the proposed data-driven algorithm can be used to determine how new patients should be assigned to the future available time slots by considering the remaining capacity. The algorithm leverages sample paths from the available offline data set to solve the single-stage stochastic program.

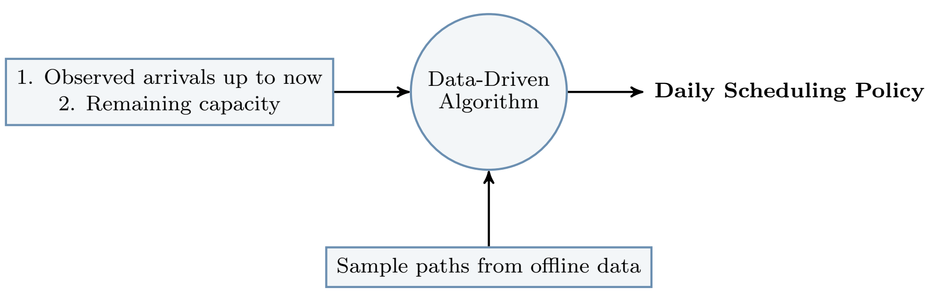

Figure 1.1: Dynamic Optimization Framework

We then extend the single-stage stochastic program to a distributionally robust optimization problem to add protection against high variability in arrivals. In the proposed robust formulation, we model the ambiguity set using a risk-averse optimization approach and propose a data-driven algorithm to solve the problem. To the best of our knowledge, this is the first paper that develops an online data-driven algorithm for a multi-class advance patient scheduling problem and proposes robust solutions against perturbations in the arrivals of patients.

Finally, we study the case of Ontario MRI services to evaluate the performance of the proposed robust formulation by comparing it with three policies including the offline policy, myopic policy, and the proposed stochastic model. We observe that the proposed robust approach outperforms the stochastic model significantly. In particular, we observe that the gap between the proposed robust and stochastic models increases as the variability in the future demand increases. This result confirms that the robust framework would be able to protect the system against the high variability in the future demand. In other words, the scheduling policy obtained based on the stochastic model cannot efficiently protect the system against the high variability in arrivals.

Second, by comparing the performance of the proposed robust approach with the offline policy, which is based on the full information about the future arrivals, we observe a relative gap of up to 75%. This shows that the proposed robust model performs well considering that the offline algorithm uses the full knowledge of the future arrivals. Furthermore, we observe a large relative gap of up to 600% between the proposed robust framework and the myopic policy. This observation indicates the negative impact of ignoring future arrivals when future time slots are assigned to patients on the performance of the system could be significant. This is mainly due to keeping low capacity for the future arrivals of high priority patients will impose a high cost on the system.

The rest of this work is organized as follows. In Chapter 2, we review the literature relevant to the advance patient scheduling problems. In Chapter 3, we develop a dynamic optimization framework for the multi-class advance patient scheduling problem and propose a dynamic stochastic method to solve it in Chapter 4. In Chapter 5, we propose a robust approach to protect the scheduling decisions against perturbations in future arrivals of patients. Using data from Ontario MRI services, we conduct a case study in Chapter 6, and conclude the paper in Chapter 7. We provide all the proofs and supplementary technical details in the Appendix.

1.1 Notation

Before proceeding, we introduce some notations. Throughout this work, all the vectors are column vectors unless otherwise specified. For any $n \in \mathbb{N}$ and $N \in \mathbb{N}$, we use $[n]$ and $]n, N]$ to denote the set $\{1, 2, \ldots, n\}$ and the set $\{n+1, n+2, \ldots, N\}$, respectively. For any column vector $x \in \mathbb{R}^n$, the ℓ_p-norm is $\|x\|_p := (\sum_{i=1}^n |x_i|^p)^{\frac{1}{p}}$. For any vector $x \in \mathbb{R}^n$ and $a \in \mathbb{R}$, we use $x \geq a$ to denote $x_i \geq a, i \in [n]$. For any $a \in \mathbb{R}$, $[a]_+ = \max\{a, 0\}$. For any matrix $A_{n\times m}$, $\vec{A}$ denotes the vectorization of A. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a measure space where Ω is a sample space equipped with the sigma algebra $\mathcal{F}$, and $\mathbb{P}$ is a measure defined on the measurable space $(\Omega, \mathcal{F})$. We denote a Banach space on this measure space by $\mathscr{L}_p(\Omega, \mathcal{F}, P)$. Let $Z = Z(\omega)$ be a random function in a Hilbert space $\mathcal{Z} = \mathscr{L}_2(\Omega, \mathcal{F}, \mathbb{P})$. For any $\zeta \in \mathcal{Z}$,

$$\langle \zeta, Z \rangle := \int_\Omega \zeta(\omega) Z(\omega) dP(\omega), \tag{1.1}$$

denotes the scalar product of Z and ζ.

Chapter 2

Literature Review

Appointment scheduling has been the focus of many papers in the literature. In general, papers in the literature of appointment scheduling can be categorized to two groups. A stream of papers consider *intra-day* scheduling where the goal is to examine how to schedule patients within a day. The main focus of the papers in this literature is usually to examine the *direct* wait time, which is defined as the time between the patient appointment time and the actual time that she receives the service, by determining the sequencing and timing of appointments within a day (see e.g. [12], [48], and [26]). There are also papers in this literature that examine the number of patients that should be scheduled in each time slot during a day when the day is divided into a number of time slots (see e.g. [33] and [62]). [18] and [2] provide a comprehensive review of the related literature.

Another stream of papers in the literature of appointment scheduling consider *inter-day* scheduling where the goal is to examine how patients should dynamically be scheduled to appointment days. This part of the literature usually focus on *indirect* wait time, which is defined as the time between when a patient requests an appointment and the time of the appointment, by considering uncertainty in the daily demand and capacity. According to [36], two main paradigms for multi-day scheduling are *allocation* scheduling and *advance* scheduling policy. Under the allocation scheduling policy, it is assumed that the service provider keeps a list of requests for the service and patients are informed on the day of their appointments. Under the advance scheduling policy, the appointment day of patients is determined when their requests are received.

Dynamic allocation scheduling is studied fairly extensively in the literature. For example, [14] considers a two-class dynamic scheduling problem and characterize an optimal policy for allocating capacity to each class of emergency and elective surgeries. [5] consider the problem of dynamically allocating a fixed capacity to several classes of patients. They partially characterize the optimal policy and develop a heuristic to allocate capacity to different classes of patients. [21] consider dynamic allocation of multiple resources to two classes of elective and emergency patients where the emergency demand should be satisfied immediately while the elective demand can wait. They prove some properties of the problem and develop several efficient policies to allocate resources to different classes. [41] also consider dynamic resource allocation to multiple classes of patients considering that patient priority is time-dependent.

Dynamic advance appointment scheduling has been also the focus of several papers in the literature. For example, [47] consider dynamic scheduling of patients with different priorities. They characterize some properties of the problem and develop heuristics based on the approximate dynamic programming to solve the problem. [34] consider dynamic patient scheduling with cancellation and no shows. They develop heuristics based on a static state-independent scheduling policy to dynamically schedule patients. [17] study multi-class dynamic resource allocation with multiple resource constraints. They develop an approximate dynamic program based on Lagrangian relaxation and constraint generation to solve the problem.

More recently, [60] considers advance patient scheduling in a system facing two classes of patients including an urgent and a regular demand class. It is assumed that while the regular demand class can be scheduled for sometime in the future, the urgent class should be served on the day of arrival. She characterizes the optimal scheduling policy and develops an effective algorithm to dynamically schedule patients. [67] extend the model of [60] to establish connections between dynamic inter-day scheduling and static intra-day scheduling. While [60] assumes that only one class of patients can be assigned to the future appointments, in this work, we allow multiple classes of patients to be scheduled to the future periods. [63] consider multi-priority patient scheduling with cancellation and develop an online scheduling algorithm to solve the problem. They show that the cost of the proposed algorithm is at most twice of the optimal offline algorithm. [53] study a multi-class multi-priority dynamic patient scheduling. They characterize some properties of the problem with deterministic service time and develop a heuristic based on approximate dynamic programming to solve the problem. [1] model and analyze advance scheduling in an asymptotic regime and characterize the asymptotic optimal policy for this problem.

A key characteristic of the advance scheduling problem is the uncertainty in the demand. Robust optimization approach has been used in many papers to tackle the uncertainties in the nature of scheduling problem. [39] propose a distributionally robust formulation for managing elective admissions in hospital and report favorable results by testing the proposed model in a simulation study based on real hospital admission data. [59] develop a robust optimisation model for solving the surgery capacity allocation problem with demand uncertainty to minimise the worst-case revenue loss. Similar to [60], they only allow one class of patients to be assigned in the future. More recently, [4] study the tactical capacity planning for an outpatient setting that was developed by [45] and present a cardinality-constrained robust method to solve the problem while protecting against uncertainty in the demand per time period.

Chapter 3

Problem Definition and Formulation

In this work, we study a *Multi-class Advance Patient Scheduling* (MAPS) problem. Consider a health unit (e.g., hospital) that receives appointment requests from patients over a *finite horizon* with T periods. We will use period and day interchangeably, since a period normally corresponds to a day in reality. We assume that patients are from I different *classes* who are distinguished based on their priorities, service times, and arrival distributions. Patients arrive sequentially over the time horizon, and they are given appointments at the time of the request. The goal is to schedule patients as their requests are received on a daily basis in the next N periods, which is called *booking horizon*, including the arrival day. This means that a patient who arrives on day $t \in [T]$, will be scheduled to receive service on day $t+n$, where $n \in \{0, \ldots, N-1\}$.

We assume that, at the time a patient request is received, the service provider assigns an appointment time slot to the patient and gives them the date as well as the exact time when their services will start. Thus, the service times can be considered deterministic, as indicated by [70]. The service time for the patient of class i is denoted by μ_i. We also assume that the service provider has limited daily regular capacity c_t units of time on day t. The regular capacity c_t is known at the beginning of day 1 for all $t \in [T+N-1]$, i.e., we do not consider any no-show in staff scheduling.

At the beginning of each period t, the service provider first observes the arrivals of day t, and then schedules them in the booking horizon before the demand realization of the future periods. Let $d_t = (d_{1,t}, d_{2,t}, \ldots, d_{I,t})$ denote the arrival vector of day t where $d_{i,t}$ is the number of patients from class i arrived on day t. We assume that the first and second moments of the patient arrival process for each class of patients are bounded. By capital D_τ, we denote the arrival vector of day $\tau \in]t, T]$ that is unknown on day t to distinguish it

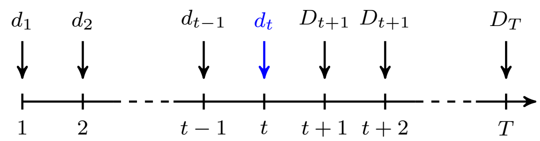

Figure 3.1: Arrival vectors on day t

from the demand realized so far in the horizon. We define the vectors $d_{[t]} = (d_1, d_2, \ldots, d_t)$ and $D_{]t,T]} = (D_{t+1}, D_{t+2}, \ldots, D_T)$ to denote the history of demand process up to day t and the vector of unknown future arrivals on day t, respectively (see Figure 3.1). We denote the joint probability distribution of $D_{]t,T]}$ by $\mathbb{P}_t$ for all $\tau \in]t, T]$. We assume that the patient arrival process is bounded meaning that the maximum daily arrival is bounded. In reality, this is a mild assumption since we can assume the arrival process is bounded by the maximum arrival in the offline data. Let Ω_t be the sample space of all random variables on day t. We assume that the arrival process is independent over periods, i.e., the arrival of day t_1 is independent of the arrival of day t_2 where $t_1 \neq t_2$.

To provide a dynamic optimization formulation, we define the action and state space of the problem as follows. Let $A_t = [a^t_{i,n}]_{I \times N}$ be the action matrix on day t, where $a^t_{i,n}$ be the number of patients of class i whose requests for the service arrived on day t and received an appointment for day $t+n$. Note that the decision matrix A_t, determined on day t, may depend on the previous actions and the arrivals up to day t, but it does not depend on the future arrivals. We formulate this problem as a multi-period (multi-stage) decision making problem over the time horizon $[0, T]$ (See Figure 3.2). Considering the action matrix A_t, the remaining capacity on day $t+n$, denoted by q^t_{n+t}, at the beginning of day t can be obtained by

$$q^t_{n+t} = q^{t-1}_{n+t} - \sum_{i=1}^{I} \mu_i . a^{t-1}_{i,n+1}, \qquad n \in \{0, \ldots, N-2\},\ t \in \{2, 3, \ldots, T\}, \tag{3.1}$$

where $q^1_k = c_k$ for all $k \in [N-1]$ since it is assumed that the system is initially empty. Then, the state of the system at beginning of day τ can be captured by

$$s_\tau = \{(d_\tau, q^\tau_\tau, \ldots, q^\tau_{\tau+N-2}) | q^\tau_{n+t} = q^{\tau-1}_{n+\tau} - \sum_{i=1}^{I} \mu_i . a^{\tau-1}_{i,n+1} \text{ for } n \in \{0, \ldots, N-2\}\}, \tag{3.2}$$

for all $\tau \in [T]$ which includes the demand of day τ and the remaining capacity of the next $N-1$ days. Note that, on day τ, we only need to include the next $N-1$ days in s_τ since patients arrived before day τ could have been scheduled for up to day $\tau + N - 2$.

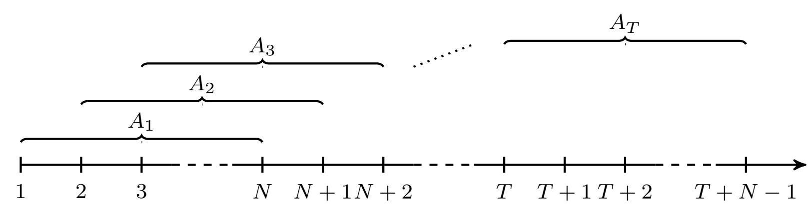

Figure 3.2: Multi-period (multi-stage) advance patient scheduling - A_t is determined on day t

Considering that all patients should receive an appointment upon the arrival of their requests, on each day t, we have the constraints $\sum_{n=0}^{N-1} a_{i,n}^t = d_{i,t}$ for all $i \in [I]$. These constraints are *deterministic* since the demand d_t is observed at the beginning of day t. We assume that the service provider may use overtime capacity and schedule more than the regular capacity c_t, i.e., we allow q_τ^t to be negative. We define a nonnegative auxiliary variable o_t to capture the overtime capacity committed on day t, where it can be determined by $o_t = \left[\sum_{i=1}^{I} \mu_i.a_{i,0}^t - q_t^t\right]_+$. Therefore, we have the capacity constraints

$$\sum_{i=1}^{I} \mu_i.a_{i,0}^t - o_t \le q_t^t, \qquad t \in [T], \tag{3.3}$$

which ensures the total appointment slots allocated to day t cannot exceed c_t unless some patients are served overtime, i.e., overtime capacity is used. We denote the decision vector of day t by $x_t = (\vec{A}_t, o_t)$. Note that adding the auxiliary variable o_t to the decision vector helps us to have a Lipschitz smooth cost function. We call the decision process $x = (x_1, x_2, \ldots, x_T)$ a *policy*. We define

$$\mathcal{C}_t(s_t) = \{x_t \mid \sum_{i=1}^{I} \mu_i.a_{i,0}^t - o_t \le q_t^t, \sum_{n=0}^{N-1} a_{i,n}^t = d_{i,t} \text{ for } i \in [I], x_t \ge 0\} \tag{3.4}$$

to denote the mapping of the state s_t to the feasible action space (set of feasible decisions) on day t. A decision is said to be feasible if $x_t \in \mathcal{C}_t(s_t)$, $t \in [T]$.

We assume that the system incurs cost for the patients waiting to being served as well as the overtime capacity utilized on each day to serve patients. Specifically, similar to [60], we assume that the system incurs the waiting cost γ_i^w per day for each patient of class i waiting to be served. We denote the overtime cost of day t by $u_t(o_t)$ and assume that $u_t(o_t)$ is a smooth convex function of o_t. The convexity of the overtime function is a common assumption in the literature (see e.g. [60]) and is compatible with reality that the service provider should pay a higher rate for each hour of overtime work. It also captures the negative consequences of the excessive amount of overtime in healthcare applications. For instance, [64] argue that overtime working hours could increase the health risk of healthcare workers and decrease productivity. Furthermore, the convex overtime cost enables the model to consider the overtime capacity in the case of emergency patients. We assume the overtime usage of resources is bounded.

The goal of the problem is to minimize the total waiting and overtime cost of scheduling over the time horizon $[0, T]$. Given decision x, the total cost can be expressed as

$$F(x) = \sum_{t=1}^{T} \left(\sum_{i=1}^{I} \gamma_i^w \sum_{n=0}^{N-1} n a_{i,n}^t + u_t(o_t) \right), \tag{3.5}$$

where the first and second terms show the linear waiting cost and the overtime cost, respectively. We ignore the cost after the time horizon since we assume $T \gg N$. Let $f(x_t)$ denote the instantaneous cost of period t given the decision of day t, x_t. Then, $f(x_t)$ can be written as,

$$f(x_t) = \sum_{i=1}^{I} \gamma_i^w \sum_{n=0}^{N-1} n a_{i,n}^t + u_t(o_t), \tag{3.6}$$

and we have $F(x) = \sum_{t=1}^{T} f(x_t)$.

We seek to design an easy-to-implement data-driven algorithm to find a feasible solution that minimizes the total cost, $F(x)$. In what follows, we first provide an offline model of the MAPS problem. We then provide two dynamic models for the problem and discuss the challenges of solving each of them.

Offline Model. An offline problem optimizes a deterministic version of the problem by assuming that the entire input sequence over the time horizon is known in advance. More specifically, in the offline setting, we assume that all arrivals $d = (d_1, d_2 \ldots, d_T)$ are known at the beginning of the horizon, i.e., at the beginning of day 1. Then, the offline problem can be formulated as

$$\begin{aligned} \mathrm{OPT}(d) = \min_{x} \quad & F(x) && && \text{(3.7a)} \\ \text{s.t.} \quad & x_\tau \in \mathcal{C}_\tau(s_\tau), && \tau \in [T], && \text{(3.7b)} \\ & (3.2), && \tau \in [T], && \text{(3.7c)} \end{aligned}$$

where it can be easily solved with the state-of-the-art solvers. In the offline problem, thanks to the prior information, we can schedule all patients over the time horizon at the beginning of day 1. However, in the main problem that future arrivals are unknown, we need to keep some capacity for the future arrivals while we schedule patients as their demand is received by the provider over time.

Dynamic Programming Model. To model the MAPS problem, one can adopt the dynamic programming approach. The value function of the problem can be stated as

$$J_t(s_t) = \min_{x_t \in \mathcal{C}_t(s_t)} f(x_t) + \mathbb{E}\left[J_{t+1}(s_{t+1}) \mid s_t\right], \tag{3.8}$$

where $\mathbb{E}[\cdot \mid s_t]$ denotes the conditional expectation. A decision x^* is optimal if and only if

$$x_t^* \in \operatorname*{arg\,min}_{x_t \in \mathcal{C}_t(s_t)} \ f(x_t) + \mathbb{E}\left[J_{t+1}(s_{t+1}) \mid s_t\right], \tag{3.9}$$

for all $t \in [T]$. Characterizing the optimal policy for the MAPS problem using the standard dynamic programming model given in (3.9) is not straightforward due to the curse of dimensionality, as indicated in the literature (see e.g., [47], [60], [11]).

Multi-Stage Stochastic Model. Another potential approach to model the MAPS problem is multi-stage stochastic programming. In each period t, d_t is revealed and the schedule x_t ideally should be the solution to the following multistage stochastic program:

$$\min_{x_t \in \mathcal{C}_t(s_t)} f(x_t) + \mathbb{E}\left[\min_{x_{t+1} \in \mathcal{C}_{t+1}(s_{t+1})} f(x_{t+1}) + \mathbb{E}\left[\cdots + \mathbb{E}\left[\min_{x_T \in \mathcal{C}_T(s_T)} f(x_T)\right]\right]\right]. \tag{3.10}$$

Note that in the multi-stage stochastic model given above, $T - t$ is the number of stages of the problem. [54] shows how the sample average approximation method, which is a promising method in two-stage stochastic problems, is not applicable for a T-*stage* stochastic problem with $T > 5$, and the number of scenarios needed to solve the problem increases *exponentially* with the number of periods T. In reality, the time horizon is usually at least a month, $T = 30$. Therefore, this problem is considered as a large-scale multi-stage stochastic program and cannot be solved with common approaches such as the sample average approximation.

Chapter 4

Dynamic Stochastic Patient Scheduling

In this section, we present an online algorithm to overcome the curse of dimensionality and solve the MAPS problem, discussed in Chapter 3. The algorithm is motivated by the idea behind *online anticipatory algorithms*. The idea behind these algorithms is to relax the *non-anticipativity* constraints and then dynamically make the decisions by solving *deterministic* optimization problems that consider all possible realizations of the random variables in the future. Although online anticipatory algorithms are sub-optimal, they perform surprisingly well ([20]). It has been shown that the online anticipatory algorithms outperform the oblivious algorithms that ignore future events.

Motivated by the online anticipatory algorithms, we propose an algorithm to solve MAPS. Specifically, similar to the online anticipatory algorithms, we relax the non anticipativity constraints, but unlike these algorithms, we solve a single stochastic optimization problem in each period, instead of multiple deterministic optimization problems. This means that we approximate problem (3.10) in two steps. On each day, we first relax the non-anticipatory constraints and then solve a single-stage stochastic optimization by *reformulating* the feasible set. Intuitively, we *schedule* patients arrived on day t while *reserving* capacity for patients who will arrive after day t by considering the unknown distribution of the arrivals in future periods. However, we do not force the problem solved in the next period to consider the capacity reserved in the current period for future arrivals. This means that we allow the optimization problem solved in each period to use all the capacity available in the future period. Note that reserving capacity for future arrivals would ensure that patients with high priority (i.e. high waiting costs) can be served relatively quickly. Also, reallocating capacity kept for future arrivals would allow the algorithm to utilize the most recent information available by observing the realized demand in each period and improve the capacity reservation.

By leveraging the above idea, we present an Online Stochastic Algorithm (OSA) for solving the MAPS problem. The outline of OSA is presented in Algorithm 1. OSA consists of two major steps. On each day t, we first observe the arrivals of day t, d_t, and update the remaining capacity of $N-1$ following days with respect to the schedule of day $t-1$, x_{t-1}. Then, for a given state s_t, we solve the following problem, which is a single-stage stochastic optimization obtained by relaxing the non-anticipatory constraints:

$$\mathrm{SDSA}(s_t) = \min_{X_t=(x_t,\bar{x}_{]t,T]})} \quad f(x_t) + \mathbb{E}\left[\sum_{\tau=t+1}^{T} f(\bar{x}_\tau)\right] \tag{4.1a}$$

$$\text{s.t.} \qquad x_t \in \mathcal{C}_t(s_t), \tag{4.1b}$$

$$\bar{x}_\tau \in \bar{\mathcal{C}}_\tau(\bar{s}_\tau), \qquad \tau \in \{t+1,\ldots,T\}. \tag{4.1c}$$

We define $\bar{x}_{]t,T]} = (\bar{x}_{t+1}, \bar{x}_{t+2}, \ldots, \bar{x}_T)$ where $\bar{x}_\tau = (\vec{A}_\tau, o_\tau)$ and $A_\tau = [a^\tau_{i,n}]_{I\times N}$. In fact, $\bar{x}_{]t,T]}$ can be interpreted as the scheduling decisions related to patients who arrive after day t, which is considered as the capacity reserved for these patients in advance. In other words, $\mu_i a^\tau_{i,n}$ is the capacity reserved on day t for patients in class i who will arrive on day τ and get an appointment for day $\tau+n$. Thus, $X_t = (x_t, \bar{x}_{]t,T]})$ denotes both the schedule and the capacity reservation determined on day t. Furthermore, $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$ captures the mapping of the state s_t to the feasible action space on day t considering the future demand. Recall that, based on the definition given in (3.4), $\mathcal{C}_\tau(s_\tau)$ denotes the mapping of the state s_t to the feasible action space assuming that the demand is known. Thus, roughly speaking, $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$ can be considered as the probabilistic version of $\mathcal{C}_\tau(s_\tau)$ since future arrivals are stochastic variables. In Section 4.1, we provide some insights on the unknown future arrivals and uncertain constraints to precisely define $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$, and then in Section 4.2, we propose a single-day stochastic algorithm (SDSA) that finds the optimal solution of problem (4.1) for a given state s_t. Later, we analyze the performance of the OSA in Section 4.3.

4.1 Unknown Future Arrivals

Since, in any period, the future demand is unknown, one challenge is how to model the constraints that include future demand in the stochastic formulation given in (4.1). The goal of this section is to address this challenge and provide a rigorous definition for $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$.

To model the scheduling of future arrivals in the stochastic model, one approach is to consider *almost surely* constraints, meaning that the model should guarantee that almost

Algorithm 1: Online Stochastic Algorithm (OSA)

Input: T, N, I, $(c_t)_{t\in[T]}$, $(\mu_i)_{i\in[I]}$, $(\gamma_i^w)_{i\in[I]}$, $(u_t(.))_{t\in[T]}$.

for $t = 1, \ldots, T$ **do**

  Step 1: Observe d_t, and

    if $t = 1$ **then** $q_k^1 = c_k, \quad \forall k \in [N-1]$,

    else $q_{n+t}^t = q_{n+t}^{t-1} - \sum_{i=1}^{I} \mu_i . a_{i,n+1}^{t-1}, \quad \forall n \in \{0, \ldots, N-2\}$.

  Step 2: Set $s_t = (d_t, q_t^t, \ldots, q_{t+N-2}^t)$, and $X_t = \text{SDSA}(s_t)$.

Output: $x^{alg} = (x_1, x_2, \cdots, x_T)$.

all future arrivals will be scheduled. There are two main criticisms to the almost surely constraints: (1) almost surely constraints treat all stochastic variables in the same way while their probability may be different, and (2) the inequalities $\sum_{n=0}^{N-1} a_{i,n}^{\tau} \geq D_{i,\tau}$ for all $\tau \in]t, T]$ and $i \in [I]$ should be satisfied for all possible realization of demand $D_{]t,T]}$ to ensure there is enough capacity to schedule the patients requesting appointments. This can be quite restrictive if the sample space of arrivals Ω_t is large since we need to take into account the maximum possible demand. This means that we will reserve a large portion of the daily capacity for future arrivals of patients with high waiting costs. This approach would not result in an efficient scheduling policy due to the high variability of daily arrivals. In particular, it leads to a high total waiting cost since it delays the scheduling of some patients to keep (reserve) capacity for the maximum demand realization of some other classes of patients.

To reduce the conservatism caused by almost surely constraints, we relax the assumption of satisfying all possible realizations of demand. Note that we only relax the constraints related to the future arrivals that determine the capacity reserved for them. Thus, the final solution obtained based on the proposed approach remains feasible and serves all patients whose requests are received. On each day, there is a deterministic constraint that guarantees to assign time slots to all patients after their requests are received by the service provider, i.e., after the demand realization. To capture the gap between the future demand and capacity reserved, we define a stochastic error for the demand of patients of class i on day τ as

$$e_{i\tau} = D_{i,\tau} - \sum_{n=0}^{N-1} a_{i,n}^{\tau}. \tag{4.2}$$

The goal is to include a constraint in the stochastic model so that the gap between the realization of the future arrivals and the capacity reserved to serve them is bounded. Thus, we bound the expected value of the squared error so that $\mathbb{E}\left[e_{i\tau}^2\right] \leq m_i$ where m_i is a positive constant. Considering this constraint, $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$ can be expressed as

$$\bar{\mathcal{C}}_\tau(\bar{s}_\tau) = \{\bar{x}_\tau \mid \sum_{i=1}^{I} \mu_i.a_{i,0}^\tau - o_\tau \leq q_\tau^\tau,\, \mathbb{E}\left[e_{i\tau}^2\right] \leq m_i \forall i \in [I],\, \bar{x}_\tau \geq 0\}. \tag{4.3}$$

In Section 6, we numerically show that considering the above constraint to reserve capacity for future arrivals performs surprisingly well. In the following theorem, we show that parameter m_i can be tuned so that an almost surely constraint on the demand satisfaction is met.

Theorem 1 (Expectation Generality) *If the expected squared error for the demand of patients in class i on day τ is bounded by a positive m_i, $\mathbb{E}\left[e_{i\tau}^2\right] \leq m_i$, for a given $\epsilon \in (0,1)$, it is guaranteed that the probabilistic constraint $Pr\left\{e_{i\tau} \leq r_i\right\} \geq 1-\epsilon$ with $r_i = \sqrt{\frac{m_i}{\epsilon}}$ is satisfied.*

Theorem 1 demonstrates that any probabilistic constraint on the demand satisfaction can be captured by a corresponding bound on the expected squared error. In other words, the parameter m_i can be tuned so that any almost surely constraint is satisfied. For instance, to ensure that the probability of ignoring more than r_i patients of class i in the future capacity reservation is at most $\epsilon \in (0,1)$, it is sufficient to include $\mathbb{E}\left[e_{i\tau}^2\right] \leq \epsilon r_i^2$ in the model.

4.2 Single-Day Stochastic Scheduling Problem

Considering the feasible action space of the problem defined in (3.4) and (4.3), the formulation of the problem provided in (4.1) can be considered as a standard stochastic optimization problem. In this section, we aim to simplify the problem by applying the Lagrangian duality approach. Let $\mathcal{Y}_t(s_t)$ denote the feasible set of problem (4.1) obtained after relaxing the constraints $\mathbb{E}\left[e_{i\tau}^2\right] \leq m_i$ for all $\tau \in]t,T]$ and $i \in [I]$. Then,

$$\mathcal{Y}_t(s_t) = \{X_t | x_t \in \mathcal{C}_t(s_t), \sum_{i=1}^{I} \mu_i.a_{i,0}^\tau - o_\tau \leq q_\tau^\tau, (3.2), \bar{x}_\tau \geq 0, \forall \tau \in]t,T]\}. \tag{4.4}$$

Applying Lagrangian approach to the stochastic problem defined in (4.1) with the constraints $\mathbb{E}\left[e_{i\tau}^2\right] \leq m_i$ is relaxed, we get the following dual problem,

$$\max_{\lambda^t \geq 0} \min_{X_t \in \mathcal{Y}_t} \quad \mathbb{E}_{D_{]t,T]}}\left[L_t(X_t, D_{]t,T]}, \lambda^t)\right], \tag{4.5}$$

where the Lagrangian L_t is defined as

$$L_t(X_t, D_{]t,T]}, \lambda^t) = f(x_t) + \sum_{\tau=t+1}^{T} \left(f(\bar{x}_\tau) + \sum_{i=1}^{I} \lambda_{i\tau}^t \left(e_{i\tau}^2 - m_i \right) \right). \tag{4.6}$$

In (4.6), $\lambda^t = (\lambda_{i\tau}^t)_{i\in[I],\tau\in]t,T]}$ denotes the dual multiplier vector. The result of applying standard Lagrangian duality theory to the problem given in (4.1) is summarized in the following theorem.

Theorem 2 *Let p^* and d^* denote the optimal solutions of (4.1) and (4.5), respectively. Also, let m_i be greater than the maximum variance of arrivals of class i,*

$$m_i > \max\{\mathrm{Var}(D_{i,t+1}), \ldots, \mathrm{Var}(D_{i,T})\}.$$

Then, equality $p^* = d^*$ holds, i.e., the duality gap between (4.1) and (4.5) is zero.

Theorem 2 shows that strong duality holds for the problem defined in (4.5), i.e., problem (4.5) can be obtained by exchanging the infimum and supremum operators of problem (4.1) as a saddle-point problem. Therefore, we can use the standard subgradient method to find the (near-)optimal solution of problem (4.5). To do so, we need to solve the sub-problem $\min\limits_{X_t\in\mathcal{Y}_t(s_t)} \mathbb{E}_{D_{]t,T]}} \left[L_t(X_t, D_{]t,T]}, \lambda^t)\right]$ for a given dual multiplier vector λ_t in each iteration of the subgradient method. Intuitively, λ^t can be interpreted as the quadratic rejection cost of patients since the term $\sum_{i=1}^{I} \lambda_{i\tau}^t m_i$ is fixed, and $\lambda_{i\tau}^t$ is the only multiplier of $e_{i\tau}^2$ in $L_t(X_t, D_{]t,T]}, \lambda^t)$. In other words, the problem $\min\limits_{X_t\in\mathcal{Y}_t(s_t)} \mathbb{E}_{D_{]t,T]}} \left[L_t(X_t, D_{]t,T]}, \lambda^t)\right]$ is a MAPS problem while we are allowed to reject patients with a convex cost function.

With a slight abuse of notation, we define $\ell_t : \mathbb{R}^{r_t} \times \Omega_t \to \mathbb{R}$ such that for a given λ^t, $\ell_t(X_t, D_{]t,T]}) = L_t(X_t, D_{]t,T]}, \lambda^t)$. Note that λ^t is often given in our analysis since it is given in the subgradient method, though we update the λ^t at the end of each iteration. Hereafter, our focus is to solve the following problem

$$\min_{X_t\in\mathcal{Y}_t(s_t)} \quad \mathbb{E}_{D_{]t,T]}} \left[\ell_t(X_t, D_{]t,T]})\right], \tag{4.7}$$

which is a standard stochastic optimization problem. The following lemma shows the convexity of both the objective function and the feasible set of problem (4.7).

Lemma 1 *The optimization problem (4.7) is a convex optimization problem, i.e., the objective function is a convex function and the feasible set is a convex set.*

Lemma 1 states that (4.7) is a stochastic convex optimization problem, that is a well-studied problem in the literature. Following the seminal work of [50], various stochastic approximation (SA) algorithms have been developed for solving this class of problems. [29] presents the first universally optimal algorithm for solving stochastic convex optimization problems (see Chapter 4 in [30] for more recent studies). We employ a unified optimal method, namely, the accelerated stochastic approximation algorithm (proposed by [15]) to optimally solve the stochastic sub-problem (4.7). In particular, this problem is a constrained convex stochastic optimization, and this method is a modified version of Nesterov's algorithm (see [44] and [43]).

Considering Theorem 2 and Lemma 1, we provide a Single-Day Stochastic Algorithm (SDSA) to solve (4.7). The outline of SDSA is presented in Algorithm 2. In particular, we utilize a subgradient method in which in each iteration, we optimally solve the stochastic problem (4.7) by implementing the accelerated stochastic approximation algorithm. The proposed algorithm is data-riven in the sense that in each iteration, we use the offline data set to randomly pick a sample arrival for each estimation of the algorithm. We refer the reader to Appendix B for more technical details.

Note that we could also tackle problem (4.1) using other approaches. For instance, this problem can be written in the form of

$$\min_{X_t \in \mathcal{Y}_t(s_t)} \quad f(x_t) + \mathbb{E}\left[\sum_{\tau=t+1}^{T} f(\bar{x}_\tau)\right] \tag{4.8a}$$

$$\text{s.t.} \quad \mathbb{E}\left[e_{i\tau}^2\right] \le m_i, \quad \forall i \in [I], \tau \in \{t+1, \dots, T\}, \tag{4.8b}$$

which is a stochastic optimization with expectation constraints. There are several recent papers in the literature that propose algorithms for solving a stochastic optimization problem with expectation constraints. For example, [31] present a stochastic approximation method for the special case of a single expectation constraint. [7] extend this approach to the case of multiple expectation constraints. Due to the convex structure of the objective function and the feasible set, these algorithms can provide the optimal solution for problem (4.8).

In Chapter 6, we use the Lagrangian duality approach to solve the problem numerically to be consistent with the robust model of MAPS problem derived in Chapter 5 based on the robust version of problem (4.7).

4.3 Performance Measure

In this section, we examine the performance of the proposed OSA by comparing it with the offline policy in the expected sense. In particular, we compare the expected cost of the proposed algorithm with the expected cost of a *clairvoyant optimal solution* that has prior knowledge of the arrivals of patients. We define the *accumulated regret* (AR) as the difference between the accumulated cost of the proposed policy and the minimum cost that could be obtained had all arrival been known in advance. Let x^{alg} denote the scheduling decision obtained based on an algorithm. The AR of the algorithm over the time horizon T is expressed as

$$\mathrm{AR}_{alg}^{T} = \ \mathbb{E}\left[F(x^{alg})\right] - \mathbb{E}\left[\mathrm{OPT}D)\right], \tag{4.9}$$

where $\mathrm{OPT}(D)$ is defined in (3.7) and $\mathbb{E}\left[\mathrm{OPT}(D)\right]$ is the expected cost of clairvoyant. We prove that AR can be decomposed into T bounded local regrets (LR). To derive LR, we define two offline problems as follows,

$$\begin{aligned} V_1^t(s_t, x_t^{alg}, D) = \min_{x_t, \dots, x_T} \quad & \sum_{\tau=t}^{T} f(x_\tau) && && (4.10a) \\ \text{s.t.} \quad & x_\tau \in \mathcal{C}_\tau(s_\tau), && \tau \in]t, T], && (4.10b) \\ & (3.2), && \tau \in]t, T], && (4.10c) \\ & x_t = x_t^{alg}, && && (4.10d) \end{aligned}$$

and

$$\begin{aligned} V_2^t(s_t, D) = \min_{x_t, \dots, x_T} \quad & \sum_{\tau=t}^{T} f(x_\tau) && && (4.11a) \\ \text{s.t.} \quad & x_\tau \in \mathcal{C}_\tau(s_\tau), && \tau \in]t-1, T], && (4.11b) \\ & (3.2), && \tau \in]t, T]. && (4.11c) \end{aligned}$$

Note that V_1^t is an offline problem to schedule patients arriving on day t to day T where the state s_t is given and decisions on day t are made based on x_t^{alg} given in (4.10d). Also, V_2^t is similar to V_1^t with a difference that there is no constraint on the decisions on day t. In fact, the expected difference between V_1^t and V_2^t captures the specific amount of regret caused by the decision x_t^{alg} made on dat t. We define the LR of decision x_t^{alg} given state s_t as

$$\mathrm{LR}_t(s_t, x_t^{alg}) = \mathbb{E}\left[V_1^t(s_t, x_t^{alg}, D) - V_2^t(s_t, D) | s_t\right]. \tag{4.12}$$

Note that in $\mathbb{E}\left[\cdots | s_t\right]$, the expectation is over the uncertain arrival process $D_{t+1}, \dots, D_T$.

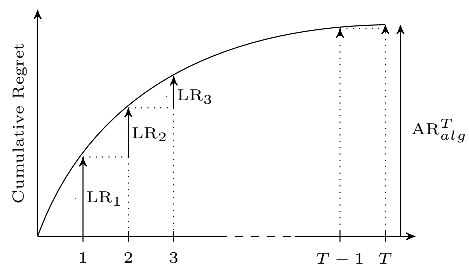

Figure 4.1: Regret Decomposition

Theorem 3 (Regret Analysis) *Let x^{alg} denote the policy obtained by Algorithm 1, and consider AR^T_{alg} and LR_t defined in (4.9) and (4.12), respectively. Then, we have* $AR^T_{alg} = \sum_{t=1}^{T} \mathbb{E}\left[LR_t(s_t, x_t^{alg})\right]$.

Theorem 3 demonstrates that AR of the proposed OSA is the sum of the local regrets, as illustrated in Figure 4.1. Intuitively, this theorem claims that the accumulated regret can be computed by adding the expected specific amount of regret caused by the decision of each day. This provides an efficient way to compute the AR of the proposed algorithm.

Chapter 5

Dynamic Robust Patient Scheduling

In the previous section, we took an adaptive approach and developed a stochastic optimization model for patient scheduling, and proposed an algorithm to solve it. In this model, patients are scheduled as their requests are received by the service provider so that the expected total waiting and overtime cost of the system is minimized. The main challenge of this approach is that it does not protect the system against the high variability in daily arrivals as it is illustrated in Figure 6.3 (see also Figure 1 in [47]). To address this challenge, in this section, we propose an Online Robust Algorithm (ORA) that protects the system against perturbations in the arrival of patients.

Similar to the online stochastic algorithm presented in Algorithm 1, we take an adaptive approach to update the remaining capacity day-to-day while, in the ORA, we utilize a robust scheduling policy named Single-Day Robust Algorithm (SDRA) instead of SDSA. Thus, the outline of the ORA is similar to Algorithm 1, but in Step 2, we solve the SDRA instead of SDSA.

More specifically, we present a distributionally robust (DR) formulation for the problem defined in (4.1) and develop an online algorithm to solve the proposed model. In Section 5.1, we propose this formulation and develop an algorithm to efficiently solve it. Using patient-level MRI data from several hospitals, in Section 6, we show that the proposed distributionally robust framework not only outperforms the adaptive stochastic approach in practice significantly but also performs close to the offline policy.

5.1 Single-Day Robust Scheduling Problem

In Chapter 4, we show that the stochastic formulation provided in (4.7) is equivalent to the problem given in (4.1) for the optimal Lagrangian multipliers. To reduce the sensitivity of the scheduling policy to high changes in the arrivals, we consider the DR formulation of problem (4.7) as follows

$$\min_{X_t \in \mathcal{Y}_t(s_t)} \max_{P_t \in \mathcal{M}(\mathbb{P}_t)} \mathbb{E}_{P_t} \left[\ell_t(X_t, D_{]tiven,T]}) \right], \tag{5.1}$$

where $\mathcal{M}_t(\mathbb{P}_t)$ denotes the ambiguity set of probability measures that capture perturbations of $\mathbb{P}_t$. In (5.1), X_t is chosen to minimize the expected cost in the worst-case perturbations. Note that (5.1) reduces to (4.7) provided that the probability measure P_t contains only the true distribution of random vector $D_{]t,T]}$. That is, the DR model is an intermediate path between the robust framework and the stochastic model (see, e.g., [49] for a recent comprehensive review of the DR optimization).

Recently, there have been many studies in the literature that propose various methods to model an ambiguity set (see e.g. [56], [13], and [42]). These methods can be mainly categorized into some specific frameworks including f-divergence-based sets, Wasserstein balls around $\mathbb{P}$, risk-aversion models, and data-driven methods (see [19]). In this paper, to rigorously define the ambiguity set $\mathcal{M}_t(\mathbb{P}_t)$, we quantify the uncertainty in arrivals by utilizing a *risk measure*, which is a mapping from a space of random variables to the extended real line. We then present a data-driven algorithm to minimize the risk measure and provide a robust scheduling policy.

More specifically, we aim to reformulate the DR formulation (5.1) to a risk measure minimization problem. [55] shows that a *coherent* risk measure can be expressed as a robust expectation with respect to a particular family of distributions (see Theorem 6.6 in [55]). Recently, [28] provide some empirical evidence suggesting the *mean-upper-semideviation* risk measure can provide protection against value function approximation errors. The mean-upper-semideviation risk measure is one of the most popular risk measures in theory and practice (see e.g. [51], [3], [10], [35], and [23]). This motivates us to model the ambiguity set $\mathcal{M}_t(\mathbb{P}_t)$ with a mean-upper-semideviation risk measure.

The general form of the mean-upper-semideviation risk measure is as follows

$$\rho(Z) = \mathbb{E}\left[Z\right] + \varkappa \left(\mathbb{E}\left[\left[Z - \mathbb{E}\left[Z\right] \right]_+^p \right] \right)^{\frac{1}{p}}, \tag{5.2}$$

where Z is a random variable, p is the order of risk measure, and the level of robustness is controlled by the parameter $\varkappa$. To adapt the risk measure (5.2) to our problem, recall that $\ell_t(X_t, D_{]t,T]})$ denotes the Lagrangian of problem (4.1) for a given dual vector λ_t. Also, $\ell_t(X_t, D_{]t,T]})$ is a function of future arrivals, thus, it is a random variable itself. Therefore, we define the random variable $Z_t = \ell_t(X_t, D_{]t,T]})$ on the measurable space $(\Omega_t, \mathcal{F}_t)$. Furthermore, choosing $p = 2$ and considering that the expectation $\mathbb{E}[Z_t]$ is bounded due to the bounded arrival process, we use the Hilbert space $\mathcal{Z}_t := \mathscr{L}_2(\Omega_t, \mathcal{F}_t, \mathbb{P}_t)$ to define the feasible space of random functions Z_t. Then, we have $Z_t \in \mathcal{Z}_t$. Therefore, the mean-upper-semideviation risk function adapted to our problem is

$$\rho(Z_t) = \mathbb{E}[Z_t] + \varkappa\sqrt{\mathbb{E}\left[[Z_t - \mathbb{E}[Z_t]]_+^2\right]}. \tag{5.3}$$

[55] shows that the risk measure in (5.3) is coherent if $\varkappa \in [0, 1]$, and a real-valued coherent risk measure has a dual representation with the form of

$$\rho(Z_t) = \max_{\zeta_t \in \mathcal{A}_t} \langle \zeta_t, Z_t \rangle = \max_{P_t : \frac{dP_t}{d\mathbb{P}_t} \in \mathcal{A}_t} \mathbb{E}_{P_t}[Z_t], \quad \forall Z_t \in \mathcal{Z}_t, \tag{5.4}$$

where $\mathcal{A}_t = \left\{\zeta_t' \in \mathcal{Z}_t \middle| \zeta_t' = 1 + \zeta_t - \mathbb{E}[\zeta_t], \|\zeta_t\|_2 \leq \varkappa, \zeta_t \leq 0\right\}$ is a non-empty convex closed set. Specifically, (5.4) shows that for any $\zeta_t \in \mathcal{A}_t$ and $Z_t \in \mathcal{Z}_t$ we can view $\langle \zeta_t, Z_t \rangle$ as an expectation taken with respect to the probability measure $\zeta_t d\mathbb{P}_t$. The expectation in (5.1) is with respect to the probability measure P_t, then, it suffices to have $dP_t = \zeta_t d\mathbb{P}_t$. Therefore, the uncertainty set $\mathcal{M}_t(\mathbb{P}_t)$ can be presented as $\mathcal{M}_t(\mathbb{P}_t) = \{P_t | \frac{dP_t}{d\mathbb{P}_t} \in \mathcal{A}_t\}$ (the same approach has been applied by [19]). By substituting (5.4) into (5.1), we get

$$\min_{X_t \in \mathcal{Y}_t(s_t)} \rho(Z_t). \tag{5.5}$$

Note that if we define the risk measure $\rho(\cdot)$ as $\rho(\cdot) := \mathbb{E}[\cdot]$ instead of (5.3), the robust formulation give in (5.5) reduces to the stochastic formulation provided in (4.7) which is a risk-neutral problem.

Recall that $Z_t = \ell_t(X_t, D_{]t,T]})$. By substituting (5.3) into (5.5), we get

$$\mathrm{SDRA}(s_t) = \min_{X_t \in \mathcal{Y}_t(s_t)} \mathbb{E}\left[\ell_t(X_t, D_{]t,T]})\right] + \varkappa\sqrt{\mathbb{E}\left[\left[\ell_t(X_t, D_{]t,T]}) - \mathbb{E}\left[\ell_t(X_t, D_{]t,T]})\right]\right]_+^2\right]}. \tag{5.6}$$

This is the single-day optimization problem that we need to solve on each day t to obtain a robust scheduling policy. Solving problem (5.6) is challenging since the objective function has a nested expectation structure. Next, we reformulate the problem (5.6) to a stochastic multi-level composition optimization problem. There are several recent papers in the literature that present algorithms for solving stochastic multi-level composition optimization problems, see, e.g., [66], [69], [52], and [6]. To develop an optimal algorithm for solving (5.6), we first characterize some properties of this optimization problem.

To this aim, first, we prove that the problem (5.6) is convex. The convexity of this problem follows from the problem's convex cost structure and the convexity of the risk measure and the feasible set. The result is summarized in the following lemma.

Lemma 2 *The robust optimization formulation given in* (5.6) *is a convex problem.*

Second, problem (5.6) can be represented in the form of a stochastic 3-level composition optimization problem as follows,

$$\min_{X_t \in \mathcal{Y}_t(s_t)} \left[H(X_t) \coloneqq h_{1,t}\Big(X_t, h_{2,t}\Big(X_t, h_{3,t}(X_t) \Big) \Big) \right], \tag{5.7}$$

where $h_{1,t}, h_{2,t} : \mathbb{R}^{r_t} \times \mathbb{R} \to \mathbb{R}$, and $h_{3,t} : \mathbb{R}^{r_t} \to \mathbb{R}$ are given by

$$h_{1,t}(X_t, \omega_2) = \mathbb{E}\left[\ell_t(X_t, D_{]t,T]}) + \varkappa \sqrt{\varepsilon^2 + \omega_2} \right], \tag{5.8a}$$

$$h_{2,t}(X_t, \omega_3) = \mathbb{E}\left[\left[\ell_t(X_t, D_{]t,T]}) - \omega_3 \right]_+^2 \right], \tag{5.8b}$$

$$h_{3,t}(X_t) = \mathbb{E}\left[\ell_t(X_t, D_{]t,T]}) \right]. \tag{5.8c}$$

In particular, converting the nested expectation structure to the above composition form enables us to benefit from the recent studies providing algorithms to solve multi-level composition functions. Note that as a minor regularization, we add the scalar ε to the second term inside the expectation in $h_{1,t}$ to ensure the Lipschitz continuity of $h_{1,t}$. Note also that the composition form of (5.7) can be written in the form of $H(X_t) = h'_{1,t} \circ h'_{2,t} \circ h'_{3,t}(X_t)$, which is a more well-known form of nested functions (see e.g. [6], [65], [22]).

Next, we show that all layer functions $h_{1,t}, h_{2,t}$ and $h_{3,t}$ are Lipschitz smooth. We summarize the results in the following Lemma.

Lemma 3 *All layer functions $h_{i,t}(\cdot)$ and their derivatives defined in the scheduling problem* (5.7) *are Lipschitz continuous.*

The above properties enable us to adopt a multi-level nested averaging stochastic gradient method proposed by [6] to solve the robust model. This method is an extension of the NASA algorithm provided in [16] (developed for a two-level composition function) to the general multi-level setting, requiring a mini-batch of sample in each iteration. They have two major assumptions on the Stochastic Oracle. First, they assume that we have access to unbiased noisy evaluations of each layer function and its gradient. Second, the outputs of stochastic oracle are independent both in each layer and between layers. These assumptions are consistent with our setting considering that we have access to an offline data set with independent samples of arrivals.

The outline of SDRA is presented in Algorithm 3. This algorithm seeks to find an approximate stationary point of the problem, which in our case can be translated to the optimal point due to the convexity of the objective function. In this algorithm, two auxiliary averaged sequences statistically estimate the gradient of the composite objective function and the inner function value by utilizing the independent sample paths from the offline data set. Then, we compute a biased estimate of a subgradient of the composite function and update the path averages accordingly. We refer the readers to Appendix C for more technical details.

Chapter 6

Case Study: Ontario MRI Services

This paper is motivated by the scheduling of MRI services in Ontario, Canada. In this section, using patient-level Ontario MRI data, we conduct an extensive computational analysis to examine the performance of the stochastic and robust models proposed in the previous sections and provide some insights into the advance scheduling problem.

A universal, publicly funded healthcare system exists in Canada. To provide access to healthcare based on need rather than the ability to pay, each provincial or territorial insurance plan includes a basket of free services at the point of care. Inequities and lengthy wait times for elective care still exist, nevertheless (see [37]). To this end, Ontario's Wait Time Strategy was unveiled in 2004 by the Ministry of Health. Beginning in 2014, MRI and CT data is being recorded to the Wait Time Information System (WTIS) to enhance the management and monitoring of MRI and CT resources.

6.1 Data Description and Numerical Setting

We have access to a part of WTIS which contains more than two million patient-level records on MRI scheduling and delivery from 2014 to 2017 in Ontario, Canada. Below, we briefly discuss the data set (for more information, see [9]).

6.1.1 Classes of Patients.

In Ontario, to receive MRI, patients need to see their physicians to specify their service type, which refers to the body part that needs to go through an MRI scan (e.g. Brain, Head

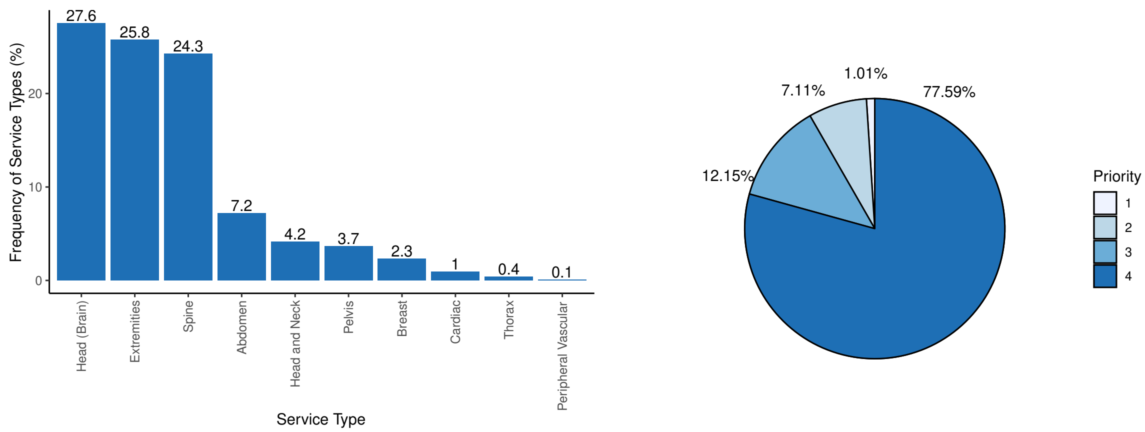

(a) (b)

Figure 6.1: Frequency of Patients' (a) Service Type and (b) Priority

and Neck, Spine, Cardiac, etc). Besides, physicians send a request to a hospital or imaging centre to set an appointment based on the available capacity. Upon receiving the request, the service provider determines the patients' priority levels using the clinical evidence and information that physicians provide. There are four priority levels: emergent (priority 1), inpatient or urgent (priority 2), semi-urgent (priority 3), and non-urgent (priority 4). By definition provided in [9], failure to diagnose and initiate treatment would result in serious morbidity/mortality for priority 1 patients, significant deterioration/deficit for priority 2 patients, moderate deterioration/deficit for priority 3 patients, and minimal deterioration/deficit for priority 4 patients. Figure 6.1 illustrates the frequency of the MRI service requests based on the service types and priorities captured in our data set. Shortly after determining the patients' priority levels, the schedulers review the clinical records and, depending on their priority and available time slots (capacity), creates an appointment for the patient. Figure 6.2 presents a general scheme of patients' journey to receive MRI service in Ontario.

Patients' priorities and the type of service determine their primary characteristics, including service duration, the arrival pattern, and the cost of delay in service. To simplify the notation, we will categorize patients into different *classes* that capture both their priority and service type simultaneously. Therefore, the number of classes is the multiplication of the number of priorities and the number of service types.

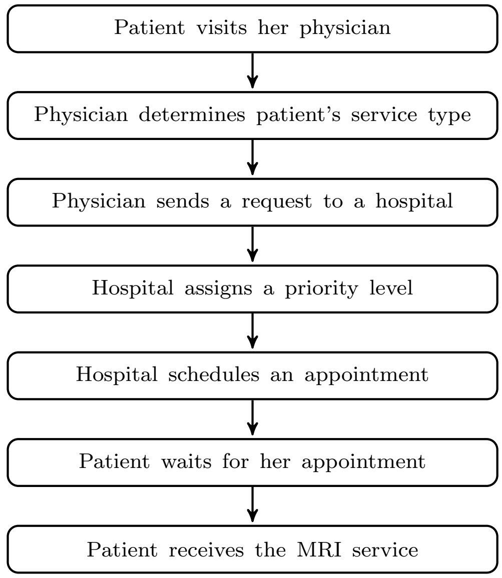

Figure 6.2: Patient's MRI Journey in Ontario

6.1.2 Service Duration.

Although the service duration is inherently random, in practice, the scheduler allocates time slots to patients based on Estimated Service Duration (ESD) at the time of making an appointment. ESD refers to the length of estimated feet in and feet out time (minutes) and not the length of scanning time (minutes) allocated for the appointment ([9]). In our data set, ESD is not recorded but the actual service duration is provided. Thus, we approximate the ESD for different types of services based on the mean of the MRI scan procedure duration, as provided in Table 6.1. The total number of observations for a descriptive summary of procedure duration, provided in the table, is 1,954,088.

6.1.3 Server's Capacity.

Server's *regular capacity* represents the number of operational hours that are anticipated for the MRI scanner during a given day. In reality, these operating hours are manually submitted ahead of time which is consistent with our setting where we assume the available capacity is given in advance. We allow the actual operational hours to exceed the regular capacity, which refers to the overtime capacity.

Service Types	ESD
Abdomen	36.25
Breast	35.02
Cardiac	59.37
Extremities	27.65
Head (Brain)	29.76
Head and Neck	38.71
Pelvis	38.90
Peripheral Vascular	37.57
Spine	27.27
Thorax	45.01

Table 6.1: Estimated Duration Scan of each service type

6.1.4 Numerical Setting.

To conduct the numerical analysis, we refer the readers to Table 6.2 for all problem parameters and numerical settings. Based on the available data, we define a baseline example and set the default parameters based on that. We will conduct a comprehensive sensitivity analysis to investigate the impact of changes in the parameters. For example, as it has been shown in Figure 6.1, there are ten service types and four priority levels for patients in the data set. Therefore, we have the arrival records for 40 classes of patients in our data set. Furthermore, we picked three random hospitals among all available hospitals in Ontario where these hospitals can be categorized as a large (L), medium (M), and small (S) size hospitals to show whether the size of the hospital can affect the performance of the scheduling policies obtained based on the proposed formulations. We assume all hospitals are initially empty at beginning of the first day in the time horizon. Figure 6.3 illustrates the daily number of service requests received by these hospitals. As shown in the figure, there is a great degree of variability and non-stationarity in the patient arrival patterns, which represents a high degree of uncertainty. Due to this high variability in the arrivals, we will show that the proposed robust formulation, which has no knowledge of the upcoming arrival pattern, outperforms the stochastic formulation significantly and performs well compared to the offline algorithm.

It remains to tune the cost parameters. We assume that the waiting cost of patients depends on their priority level. This is compatible with the real-world setting since the Ontario Ministry of Health and Long-Term Care developed wait time targets based on the priority level of patients, emphasizing the fact that the impact of waiting on patients mainly depends on their priority levels. Based on the proposed wait targets, the majority

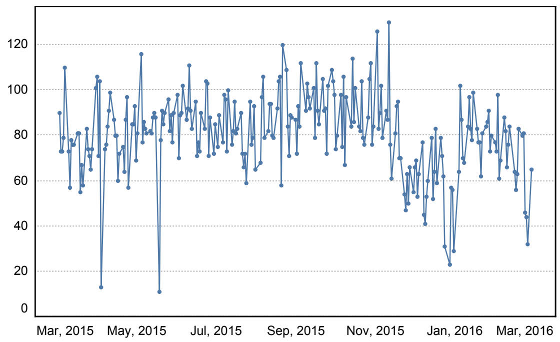

(a)

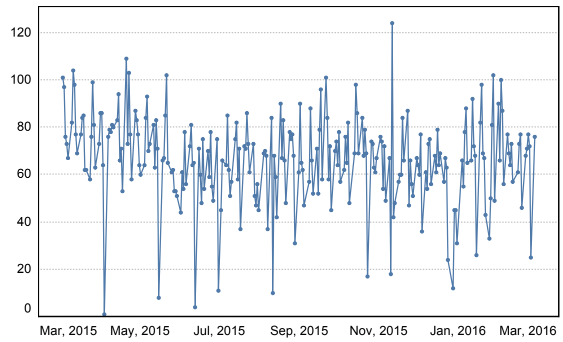

(b)

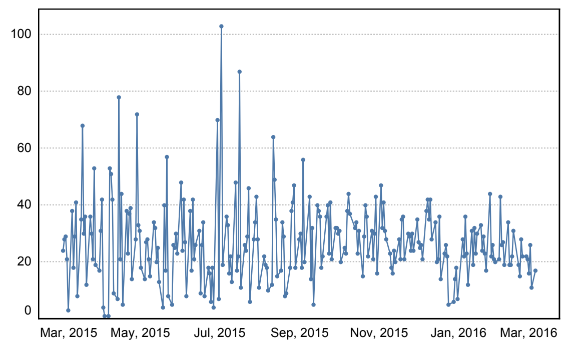

(c)

Figure 6.3: Patients Daily Arrivals - (a) Hospital L (b) Hospital M (c) Hospital S

(i.e. 90%) of patients with priority 1 should be served within one day, priority 2 within 2 days, priority 3 within 10 days, and priority 3 within 28 days. Although we do not have a hard constraint to meet the wait time targets, we can set the waiting cost parameters to consider the priority level of patients in their scheduling orders. Thus, we set the waiting cost parameters such that $28\gamma_4^w = 10\gamma_3^w = 2\gamma_2^w = \gamma_1^w$. We will show the impact of this waiting cost parameter setting on the percentage of patients served during their wait time targets.

We also need to set the regular capacity over the time horizon for the hospitals as we discussed in section 6.1.3. Our available data set only contains the arrival process of patients and we do not have access to the operating hours the hospitals submitted in advance. Therefore, we set the regular daily capacity based on the expected daily arrivals to each hospital. However, we conduct a sensitivity analysis in section 6.2 to investigate the impact of the capacity on the performance of the proposed policies.

Parameter	Notation	Value
Time horizon	T	60 days
Booking horizon	N	15 days
Number of classes	I	40
Regular Capacity for Hospital L	c	30 hours
Regular Capacity for Hospital M	c	22 hours
Regular Capacity for Hospital S	c	14 hours
Maximum overtime per day	O	4 hours
Overtime Cost Rate	γ^o	0.0075
Waiting Cost of Priority 1 Patients	γ_1^w	28
Waiting Cost of Priority 2 Patients	γ_2^w	14
Waiting Cost of Priority 3 Patients	γ_3^w	2.8
Waiting Cost of Priority 4 Patients	γ_4^w	1
Number of samples in sample pool	S	5000
Number of iteration in stochastic algorithm	R	2000
Number of iteration in robust algorithm	K	2000
Robust coefficient	$\varkappa$	0.5
Regularization epsilon	ϵ	0.1

Table 6.2: Parameters and the Setting of Algorithms

Finally, we approximate the overtime cost with a quadratic cost function which is a smooth convex function. We tune the constant rate γ^o of the quadratic overtime cost to $\gamma_1^w = 59.37^2\gamma^o$. Note that the maximum expected service duration in our data is 59.37. Thus, by setting $\gamma_1^w = 59.37^2\gamma^o$, there is a chance that some patients are scheduled on the day of their arrivals if there is no regular capacity available.

6.2 Empirical Performance

We evaluate the performance of the proposed robust formulation by comparing it with three other policies. The first benchmark is OSA proposed in Chapter 4. Second, we compare ORA with a *myopic policy* which is a naive policy and schedules patients by solving an offline problem assuming no patients arrive in the future periods. In fact, myopic policy optimally schedules patients who arrive on each day while ignoring the future arrivals. Finally, we evaluate the performance of ORA by comparing it with the *offline policy* which is based on the full information of the arrivals and provides the lowest cost. Speaking of the cost of the scheduling, we aim to examine how far the costs of ORA and OSA are from the costs of the offline policy (as a lower bound) and the myopic policy (as an upper bound).

We define the *Relative Gap* between the cost of two polices as the metric to compare the performance of the two policies.

Definition 1 (Relative Gap) *Let C_A and C_B be the expected total cost of patient scheduling based on policies A and B, respectively. Then, the relative gap between policy A and policy B is defined as:*

$$G_{A,B} = \frac{C_A - C_B}{C_B}. \tag{6.1}$$

To compute the relative gap defined in (6.1), we randomly pick 200 samples of arrivals from the data set to create a *test set*. These samples are kept fixed during all the analyses.

6.2.1 Sensitivity Analysis on Booking Horizon.

We first compare the performance of ORA and OSA as the length of the booking horizon N varies, as illustrated in Figure 6.4. We observe that $G_{\text{OSA,ORA}}$ increases as the booking horizon increases in all three hospitals. In fact, the performance of OSA and ORA is close to each other for a small booking horizon of 3 days while the relative gap increases to $64 - 110\%$ when it is possible to schedule patients for 30 days in advance. Note that as the number of periods in the booking horizon increases, the variability in the total arrivals during the booking horizon increases. Thus, Figure 6.5 shows that as the variability in the total arrivals during the booking horizon increases, the performance of the robust model compared with the stochastic model improves significantly. This observation confirms that when the variability in the demand during the booking horizon becomes high, the scheduling policy obtained based on the stochastic formulation of the problem cannot efficiently protect the system against the high variability in the future demand. However, the robust approach provides policies that take into account the possibility of observing a high demand in the future when they allocate future time slots to the patients. Thus, policies obtained based on the robust model are much more efficient which leads to a gap of up to 110% between OSA and ORA.

Note that from the risk-aversion perspective, having the term $\varkappa\sqrt{\mathbb{E}\left[[Z_t - \mathbb{E}[Z_t]]_+^2\right]}$ in the risk measure enables us to control the variation in the loss function. Therefore, the higher variability in the problem causes the larger relative gap between the robust approach and stochastic approach which highlights the importance of robust solutions in the case of high variability in arrivals.

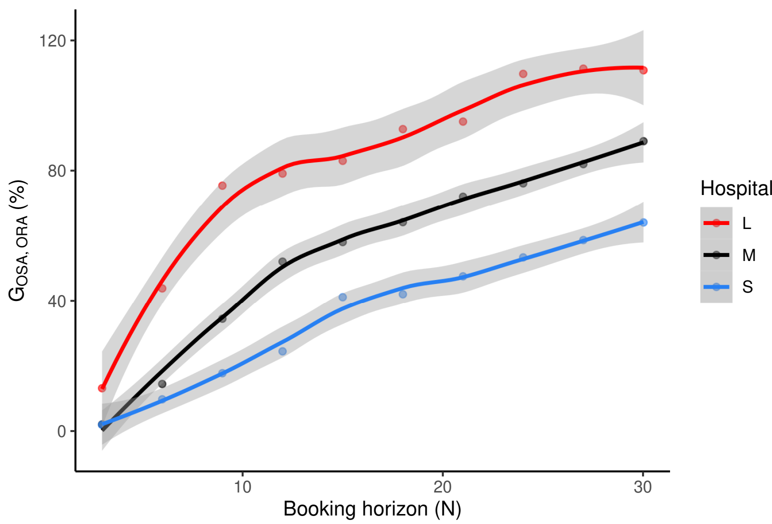

Figure 6.4: Sensitivity Analysis on Booking Horizon - Relative Gap between OSA and ORA

Based on Figure 6.4, we also observe that the gap between ORA and OSA increases as the size of the hospital increases. Note that as a hospital becomes larger the variability in the daily arrivals to the hospital increases, as shown in Table 6.3. Thus, as mentioned above, the proposed robust approach can protect the system against high variability in the arrivals during the booking horizon more efficiently than the proposed stochastic model.

Hospital	SD	IQR	MAD
L	33.76	68	21
M	29.23	53	19
S	15.25	25	12

Table 6.3: Statistical Dispersion of Daily Arrivals at Three Hospitals

To evaluate the performance of the robust approach, we compare it with the offline and myopic policies (see Figure 6.5).

To provide a simple insight on the myopic policy, assume that we have no overtime capacity. Then, the solution in myopic policy can be characterized as the well-known $c\mu$-rule scheduling policy ([61]). In other words, the agent sorts the patients based on the ratio of their waiting cost and service duration and then schedules them accordingly. Therefore, all patients will be scheduled as soon as possible assuming no patients coming in the future.

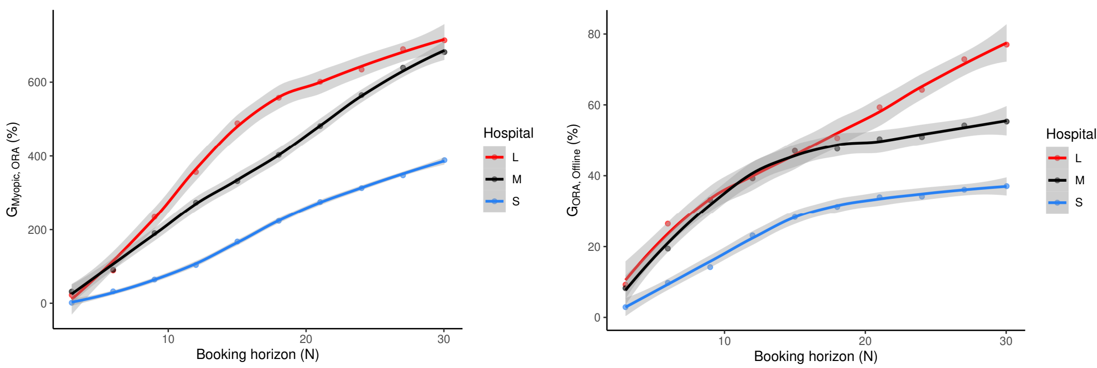

(a) Myopic Policy to ORA (b) ORA to Offline Policy

Figure 6.5: Sensitivity Analysis on Booking Horizon - Relative Gap of

The myopic policy is a good benchmark to compare with the ORA since it can indicate the effect of keeping a reserved capacity for the future arrivals of patients. Figure 6.5 (a) demonstrates that the relative gap between the myopic policy and the proposed ORA is very small when the booking horizon is one day. This is expected since both policies should schedule all arriving patients on the same day of their arrival. However, when the booking horizon increases to 30 days, the gap could go above 600%. This is mainly because the myopic policy ignores the future arrivals. Specifically, since the difference between the waiting cost of patients with different priorities is high (see Table 6.2), keeping no capacity for the future arrivals of high-priority patients will impose a high cost on the system.

We also compare the proposed ORA with the offline policy which refers to a policy obtained assuming that the arrivals of all future patients are known in advance. Note that the performance of the offline policy is not achievable in reality due to the assumption of there is no uncertainty in the arrival process. Figure 6.5 (b) shows the relative gap between ORA and the offline policy. We observe that the gap is increasing in the length of the booking horizon. This is expected since as the number of periods in the booking horizon increases, the variability in the future arrivals increases, which is known to the offline algorithm but unknown to the ORA. Our numerical results show that the gap between the offline algorithm and ORA varies between zero percent and 75% when the booking horizon increases from one day to 30 days. Considering that the offline algorithm is based on ignoring the variability completely and the full knowledge of the future arrivals, these results show that the proposed robust approach works well.

6.2.2 Sensitivity Analysis on Time Horizon.

In this section, we conduct an analysis similar to section 6.2.1 to examine the impact of the time horizon on the performance of the proposed algorithms. We fix the booking horizon to 15 days while the time horizon varies from 20 days to 60 days. Thus, the ratio N/T changes from 25% to 75%.

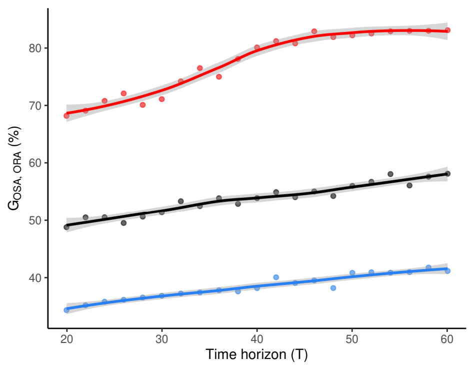

(a) OSA to ORA

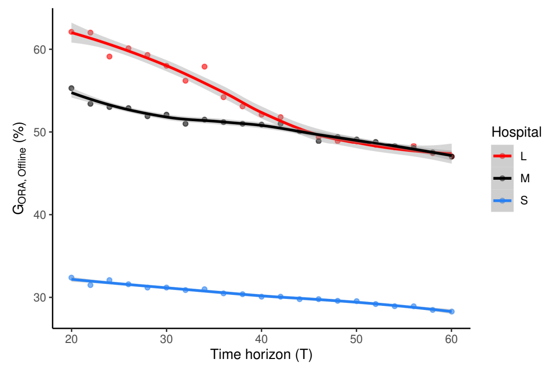

(b) ORA to Offline Policy

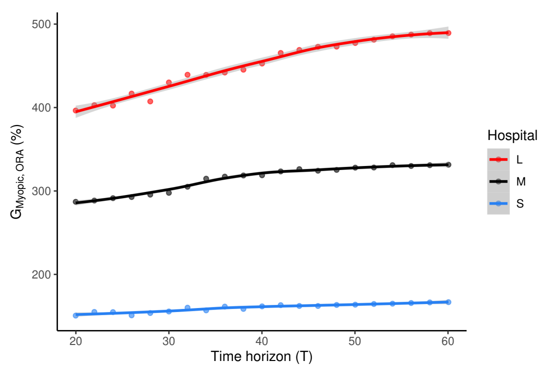

(c) Myopic Policy to ORA

Figure 6.6: Sensitivity Analysis on Time Horizon

Figure 6.6 (a) shows that as the number of periods in the time horizon increases, the gap between the proposed robust model and stochastic formulation increases. This is mainly due to an increase in the total variability observed as the time horizon increases. We observe a similar pattern in the gap between the proposed robust approach and the myopic policy, based on Figure 6.6 (c). Moreover, as illustrated in Figure 6.6 (b), the relative gap between the robust model and the offline algorithm decreases as the time horizon increases. This indicates that the performance of the proposed robust model improves as the number of periods in the time horizon increases. Recall that in the proposed approaches, we adaptively move forward through the time and schedule the patients. Thus, when the ratio N/T is high, i.e. the time horizon is short relative to the booking horizon, the effect of the end of time horizon on the overall performance of the algorithm could be significant. As the time horizon increases, while the booking horizon remains the same, the impact of the end of time horizon decreases, and thus, the impact of the robust approach becomes more significant resulting in an improvement in the performance of the robust model compared with the offline algorithm.

6.2.3 Sensitivity Analysis on Capacity.

In this section, we examine the impact of the regular capacity on the performance of the proposed ORA. To address this, we scale the regular capacity of each hospital by the factors of 0.9, 0.95, 1, 1.05, and 1.1.

Figure 6.7 (a) shows the relative gap between OSA and ORA. We observe that the relative gap is at the maximum level when there is no change in the regular capacity in all three hospitals. In other words, both increasing and decreasing the capacity leads to a smaller relative gap between OSA and ORA. This observation is consistent with our intuition since increasing capacity expands the feasible set defined in (4.4) and reduces patients' waiting time due to the presence of the waiting costs. As an extreme example, if hospitals had infinite capacity, any policy that schedules patients on the day of their arrival would lead to the lowest cost; thus, OSA, ORA, myopic, and offline policies would be optimal. In general, as the regular capacity increases, the impact of reserved capacity for future arrivals becomes less significant. As a result, the relative gap between the robust and stochastic approaches decreases. Moreover, when the regular capacity is tight, patients should wait for a long time to be served and thus the impact of the waiting cost dominates the effect of reserving capacity for future arrivals. This leads to a slight decrease in the relative gap between ORA and OSA, as illustrated in Figure 6.7 (a). We also observe that the relative gap is more sensitive to the increment of the capacity than its decrement, as depicted in Figure 6.7 (a).

Figure 6.7 (b) and (c) illustrate the relative gap between ORA and the offline and myopic policies, respectively. We first observe that increasing the capacity makes the value of efficient scheduling less significant, and thus, both $G_{\text{ORA,offline}}$ and $G_{\text{Myopic,ORA}}$ decrease. We also observe that decreasing the capacity leads to an increase in the relative gap between the ORA and the offline policy. This is mainly due to the fact that when the

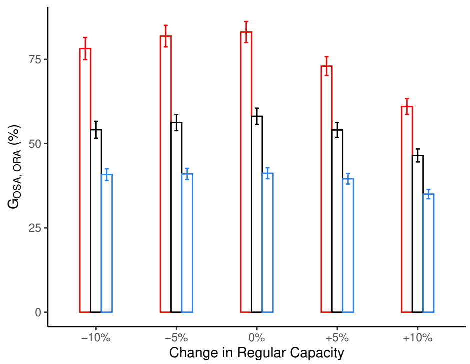

(a) OSA to ORA

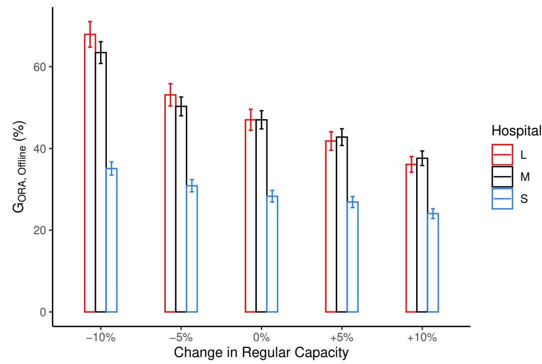

(b) ORA to Offline Policy

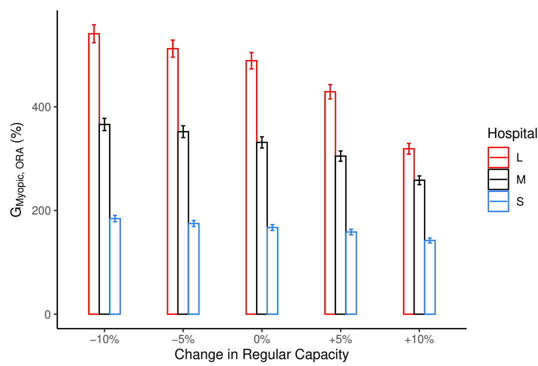

(c) Myopic Policy to ORA

Figure 6.7: Sensitivity Analysis on Regular Capacity

capacity is tight, the impact of having not enough capacity for high-priority patients would be significant since they need to wait longer compared to the case where the capacity is not tight. Finally, we observe that the large hospital is more sensitive to the change in the regular capacity in all numerical results shown in the Figure 6.7. This is due to the higher variability in the arrivals of patients at the larger hospitals.

6.2.4 Sensitivity Analysis on Robust Coefficient.

As we discussed in Chapter 5, the coefficient $\varkappa$ controls the level of robustness of the proposed ORA. In this section, we provide some computational results in Figure 6.8 to gain insights into the impact of this coefficient. First, Figure 6.8 (a) shows the ORA performs the same as OSA when $\varkappa = 0$. This is because we have $\varkappa\sqrt{\mathbb{E}\left[[Z_t - \mathbb{E}\left[Z_t\right]]_+^2\right]} = 0$ which makes the objective function of the robust model the same as the stochastic model (see (4.7) and (5.6)). Second, both Figure 6.8 (a) and (b) show the trade-off between $\mathbb{E}\left[\ell_t(X_t, D_{]t,T]})\right]$ and $\varkappa\sqrt{\mathbb{E}\left[[Z_t - \mathbb{E}\left[Z_t\right]]_+^2\right]}$. We observe that the coefficient $\varkappa$ is computationally tunable by comparing it with the OSA or the offline policy. Based on these figures, we observe that $\varkappa \approx 0.5$ would lead to the best performance of the robust model. Third, the tuned robust coefficient for the larger hospital is slightly greater than the smaller one. Intuitively, the larger robust coefficient controls the higher degree of variability in large hospitals.

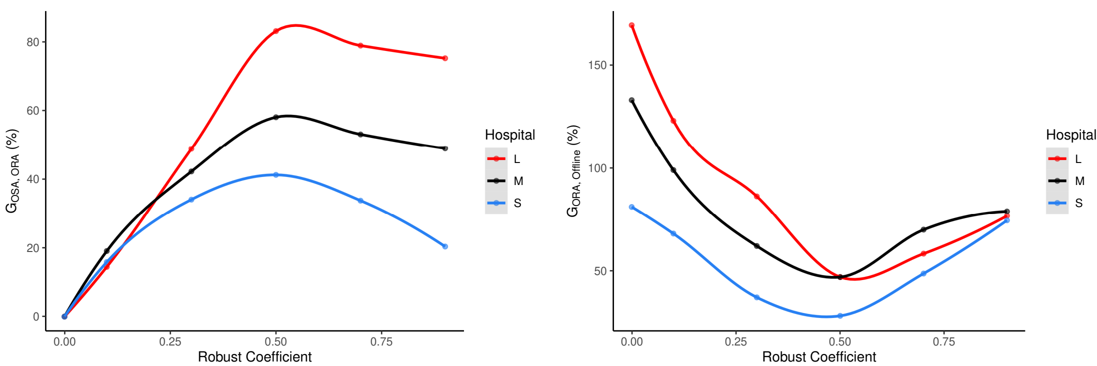

(a) OSA to ORA

(b) ORA to Offline Policy

Figure 6.8: Sensitivity Analysis on Robust Coefficient

6.2.5 Patient Waiting Time.

As we discussed in 6.1.4, in Ontario, wait time targets are defined for different priorities to ensure patients get timely access to the service. Based on the wait targets, the goal is to serve 90% of patients with priority 1, 2, 3, and 4 within one, two, ten, and twenty-eight days, respectively. To understand the waiting time of different priorities when patients are scheduled based on the proposed robust approach, we run the ORA for the 200 samples and obtain the waiting time. Figure 6.9 shows the distribution of patients' waiting time for the three hospitals. As these figures illustrate, the majority of patients are served within their wait targets at all three hospitals. Moreover, as summarized in Table 6.4, more than 90% of patients with priority 1, 2, and 4 are served within their wait targets. However, around 80% of patients of priority 3 are served in 10 days (their wait target).

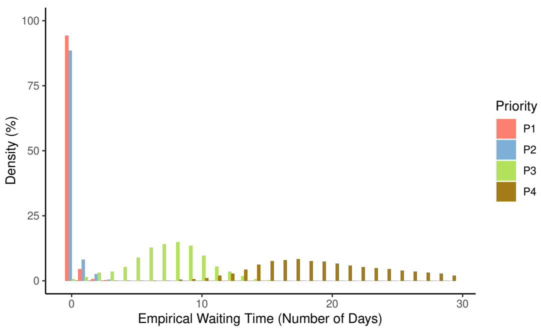

(a) Hospital L

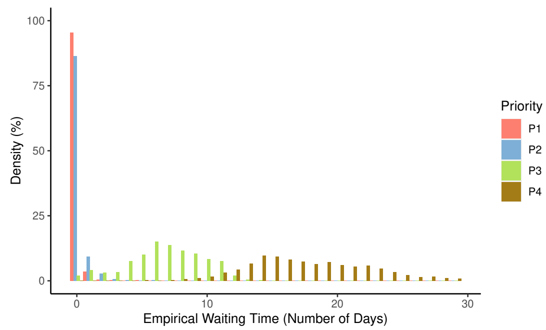

(b) Hospital M

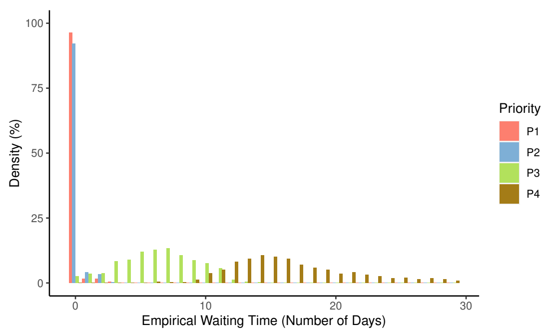

(c) Hospital S

Figure 6.9: Empirical Waiting Time Obtained by the ORA

Hospital\Priority	P1	P2	P3	P4
L	94.21	96.63	78.33	95.15
M	95.49	95.68	81.28	98.05
S	96.38	96.31	84.62	97.64

Table 6.4: Empirical Service Level (%)

Chapter 7

Conclusion and Future Research

In this work, we considered a high-dimensional multi-class advance patient scheduling problem in which patients of different classes have different service duration and arrive to the system randomly over a time horizon. It is assumed that the system incurs different waiting costs per unit of time if patients are not served immediately, which depends on the class of patients. It is known in the literature that multi-class advance dynamic patient scheduling, in general, is a challenging problem due to the high variability in the daily arrivals of patients and the high dimension of the state and action space of the problem. In this research, we presented a dynamic data-driven robust framework to tackle these challenges. Generally speaking, on each day, the proposed data-driven robust algorithm can be utilized to determine how new patients should be assigned to future available time slots considering the remaining capacity in the future.

Methodologically speaking, a novel dynamic optimization framework was developed by relaxing the non-anticipativity constraints and reformulating the constraints associated with future demand satisfaction. We approximately decomposed the multi-stage stochastic formulation of the original problem into multiple single-stage stochastic programs. Moreover, to protect the scheduling policy against the perturbations in arrivals, we presented a distributionally robust formulation in which the ambiguity set is defined based on a coherent risk measure. We characterized some properties of the problem and showed that the proposed formulation is convex. Using these properties, we developed an efficient data-driven algorithm to solve the problem.

Using Ontario MRI data, we examined the performance of the proposed models. We observed that the proposed robust policy accurately prioritizes patients from different classes and the majority of patients are served within their wait targets. We also evaluated the performance of the proposed robust formulation by comparing it with three benchmarks including a myopic policy, an offline algorithm, and a stochastic formulation of the problem. We observed that the proposed robust policy outperforms the others significantly. We also did an extensive sensitivity analysis of the problem parameters to provide insights into the advance scheduling problem. We observed that a significant improvement could be achieved, compared to the other policies, by employing the proposed robust policy, when there is a high degree of variability in the arrival process. In addition, we observed that the proposed robust model performs well compared to an offline policy, which is based on the complete knowledge of the future arrivals. These results demonstrate that the robust approach would be capable of protecting the system against the high variability in future demand and efficiently scheduling patients over the time horizon.

In this work, there are several assumptions and limitations that could be addressed in future research. First, appointment cancellations and patient no-shows are two potential directions that can be considered to extend the proposed model in this research. Moreover, we assumed that patients get an appointment for their service at the time that their requests are received by the service provider. This assumption can also be relaxed in the sense that patients would be able to wait for a short time before they receive their appointment date. In other words, the scheduler can defer the scheduling decision for a period of time to collect more information about the future arrivals and then decide how the available time slots should be allocated to the patients waiting to receive a service. In terms of methodology, we modeled the ambiguous set that captures future arrivals using a risk-averse approach, whereas different approaches exist in the literature for modeling the ambiguity set (see e.g. [56], [13], and [42]). One can examine how different ambiguity sets would affect the performance of the robust formulation. Furthermore, another interesting approach to tackle a dynamic scheduling problem is to develop an accurate prediction model for future arrivals and solve a simple optimization problem in each period using a dynamic approach. However, the challenging question would be how the performance of the scheduling policy would be affected by the error in the prediction model, which can be considered as a direction for future work.

References

[1] Hossein Abouee Mehrizi, Mohamad Kazem Shirani Faradonbeh, and Mohamad Sadegh Shirani Faradonbeh. Asymptotic analysis of multi-class advance patient scheduling. *Available at SSRN*, 2022.

[2] Amir Ahmadi-Javid, Zahra Jalali, and Kenneth J Klassen. Outpatient appointment systems in healthcare: A review of optimization studies. *European Journal of Operational Research*, 258(1):3–34, 2017.

[3] Shabbir Ahmed, Ulaş Çakmak, and Alexander Shapiro. Coherent risk measures in inventory problems. *European Journal of Operational Research*, 182(1):226–238, 2007.

[4] Nazanin Aslani, Onur Kuzgunkaya, Navneet Vidyarthi, and Daria Terekhov. A robust optimization model for tactical capacity planning in an outpatient setting. *Health Care Management Science*, 24(1):26–40, 2021.

[5] Nur Ayvaz and Woonghee Tim Huh. Allocation of hospital capacity to multiple types of patients. *Journal of Revenue and Pricing Management*, 9(5):386–398, 2010.

[6] Krishnakumar Balasubramanian, Saeed Ghadimi, and Anthony Nguyen. Stochastic multilevel composition optimization algorithms with level-independent convergence rates. *SIAM Journal on Optimization*, 32(2):519–544, 2022.

[7] Kinjal Basu and Preetam Nandy. Optimal convergence for stochastic optimization with multiple expectation constraints. *arXiv preprint arXiv:1906.03401*, 2019.

[8] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[9] . Cancer Care Ontario. Diagnostic imaging wait times and efficiencies. Data standardization guide (version 7), Ontario Health, 2018.

[10] Zhiping Chen and Yi Wang. Two-sided coherent risk measures and their application in realistic portfolio optimization. *Journal of Banking & Finance*, 32(12):2667–2673, 2008.

[11] Jiajun Dai, Na Geng, and Xiaolan Xie. Dynamic advance scheduling of outpatient appointments in a moving booking window. *European Journal of Operational Research*, 292(2):622–632, 2021.

[12] Brian Denton and Diwakar Gupta. A sequential bounding approach for optimal appointment scheduling. *IIE transactions*, 35(11):1003–1016, 2003.

[13] John Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*, 2018.

[14] Yigal Gerchak, Diwakar Gupta, and Mordechai Henig. Reservation planning for elective surgery under uncertain demand for emergency surgery. *Management Science*, 42(3):321–334, 1996.

[15] Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.

[16] Saeed Ghadimi, Andrzej Ruszczynski, and Mengdi Wang. A single timescale stochastic approximation method for nested stochastic optimization. *SIAM Journal on Optimization*, 30(1):960–979, 2020.

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

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

[19] Mert Gürbüzbalaban, Andrzej Ruszczyński, and Landi Zhu. A stochastic subgradient method for distributionally robust non-convex learning. *arXiv preprint arXiv:2006.04873*, 2020.

[20] Pascal Van Hentenryck and Russell Bent. *Online stochastic combinatorial optimization.* The MIT Press, 2006.

[21] Woonghee Tim Huh, Nan Liu, and Van-Anh Truong. Multiresource allocation scheduling in dynamic environments. *Manufacturing & Service Operations Management*, 15(2):280–291, 2013.

[22] Wei Jiang, Bokun Wang, Yibo Wang, Lijun Zhang, and Tianbao Yang. Optimal algorithms for stochastic multi-level compositional optimization. *arXiv preprint arXiv:2202.07530*, 2022.

[23] Dionysios S Kalogerias and Warren B Powell. Zeroth-order stochastic compositional algorithms for risk-aware learning. *SIAM Journal on Optimization*, 32(2):386–416, 2022.

[24] Esmaeil Keyvanshokooh, Pooyan Kazemian, Mohammad Fattahi, and Mark P Van Oyen. Coordinated and priority-based surgical care: An integrated distributionally robust stochastic optimization approach. *Production and Operations Management*, 31(4):1510–1535, 2022.

[25] Esmaeil Keyvanshokooh, Cong Shi, and Mark P Van Oyen. Online advance scheduling with overtime: A primal-dual approach. *Manufacturing & Service Operations Management*, 23(1):246–266, 2021.

[26] Qingxia Kong, Shan Li, Nan Liu, Chung-Piaw Teo, and Zhenzhen Yan. Appointment scheduling under time-dependent patient no-show behavior. *Management Science*, 66(8):3480–3500, 2020.

[27] Renata Kopach, Po-Ching DeLaurentis, Mark Lawley, Kumar Muthuraman, Leyla Ozsen, Ron Rardin, Hong Wan, Paul Intrevado, Xiuli Qu, and Deanna Willis. Effects of clinical characteristics on successful open access scheduling. *Health care management science*, 10(2):111–124, 2007.

[28] Ümit Köse and Andrzej Ruszczyński. Risk-averse learning by temporal difference methods with markov risk measures. *The Journal of Machine Learning Research*, 22(1):1800–1833, 2021.

[29] Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.

[30] Guanghui Lan. *First-order and stochastic optimization methods for machine learning.* Springer, 2020.

[31] Guanghui Lan and Zhiqiang Zhou. Algorithms for stochastic optimization with function or expectation constraints. *Computational Optimization and Applications*, 76(2):461–498, 2020.

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

[33] Nan Liu, Peter M Van De Ven, and Bo Zhang. Managing appointment booking under customer choices. *Management Science*, 65(9):4280–4298, 2019.

[34] Nan Liu, Serhan Ziya, and Vidyadhar G Kulkarni. Dynamic scheduling of outpatient appointments under patient no-shows and cancellations. *Manufacturing & Service Operations Management*, 12(2):347–364, 2010.

[35] Wann-Jiun Ma, Chanwook Oh, Yang Liu, Darinka Dentcheva, and Michael M Zavlanos. Risk-averse access point selection in wireless communication networks. *IEEE Transactions on Control of Network Systems*, 6(1):24–36, 2018.

[36] James M Magerlein and James B Martin. Surgical demand scheduling: a review. *Health services research*, 13(4):418, 1978.

[37] Danielle Martin, Ashley P Miller, Amélie Quesnel-Vallée, Nadine R Caron, Bilkis Vissandjée, and Gregory P Marchildon. Canada's universal health-care system: achieving its potential. *The Lancet*, 391(10131):1718–1735, 2018.

[38] Michael L McManus, Michael C Long, Abbot Cooper, James Mandell, Donald M Berwick, Marcello Pagano, and Eugene Litvak. Variability in surgical caseload and access to intensive care services. *The Journal of the American Society of Anesthesiologists*, 98(6):1491–1496, 2003.

[39] Fanwen Meng, Jin Qi, Meilin Zhang, James Ang, Singfat Chu, and Melvyn Sim. A robust optimization model for managing elective admission in a public hospital. *Operations research*, 63(6):1452–1467, 2015.

[40] Luc Mercier and Pascal Van Hentenryck. Performance analysis of online anticipatory algorithms for large multistage stochastic integer programs. In *IJCAI*, pages 1979–1984, 2007.

[41] Daiki Min and Yuehwern Yih. Managing a patient waiting list with time-dependent priority and adverse events. *RAIRO-Operations Research*, 48(1):53–74, 2014.

[42] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.

[43] Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.

[44] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $\mathcal{O}(1/k^2)$. In *Doklady an ussr*, volume 269, pages 543–547, 1983.

[45] Thu Ba T Nguyen, Appa Iyer Sivakumar, and Stephen C Graves. A network flow approach for tactical resource planning in outpatient clinics. *Health care management science*, 18(2):124–136, 2015.

[46] Jonathan Patrick. A markov decision model for determining optimal outpatient scheduling. *Health care management science*, 15(2):91–102, 2012.

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

[48] Jin Qi. Mitigating delays and unfairness in appointment systems. *Management Science*, 63(2):566–583, 2017.

[49] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.

[50] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

[51] R Tyrrell Rockafellar, Stan Uryasev, and Michael Zabarankin. Generalized deviations in risk analysis. *Finance and Stochastics*, 10(1):51–74, 2006.

[52] Andrzej Ruszczynski. A stochastic subgradient method for nonsmooth nonconvex multilevel composition optimization. *SIAM Journal on Control and Optimization*, 59(3):2301–2320, 2021.

[53] Antoine Sauré, Mehmet A Begen, and Jonathan Patrick. Dynamic multi-priority, multi-class patient scheduling with stochastic service times. *European Journal of Operational Research*, 280(1):254–265, 2020.

[54] Alexander Shapiro. On complexity of multistage stochastic programs. *Operations Research Letters*, 34(1):1–8, 2006.

[55] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczynski. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.

[56] Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.

[57] Clifford Stein, Van-Anh Truong, and Xinshang Wang. Advance service reservations with heterogeneous customers. *Management Science*, 66(7):2929–2950, 2020.

[58] Greg Sutherland, N Russell, Robyn Gibbard, and Alexandru Dobrescu. The value of radiology, part ii. In *The Conference Board of Canada*, 2019.

[59] Jiafu Tang and Yu Wang. An adjustable robust optimisation method for elective and emergency surgery capacity allocation with demand uncertainty. *International Journal of Production Research*, 53(24):7317–7328, 2015.

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

[61] Jan A Van Mieghem. Dynamic scheduling with convex delay costs: The generalized c— mu rule. *The Annals of Applied Probability*, pages 809–833, 1995.

[62] Shan Wang, Nan Liu, and Guohua Wan. Managing appointment-based services in the presence of walk-in customers. *Management Science*, 66(2):667–686, 2020.

[63] Xinshang Wang and Van-Anh Truong. Multi-priority online scheduling with cancellations. *Operations Research*, 66(1):104–122, 2018.

[64] Kapo Wong, Alan HS Chan, and SC Ngan. The effect of long working hours and overtime on occupational health: a meta-analysis of evidence from 1998 to 2018. *International journal of environmental research and public health*, 16(12):2102, 2019.

[65] Tesi Xiao, Krishnakumar Balasubramanian, and Saeed Ghadimi. A projection-free algorithm for constrained stochastic multi-level composition optimization. *arXiv preprint arXiv:2202.04296*, 2022.

[66] Shuoguang Yang, Mengdi Wang, and Ethan X Fang. Multilevel stochastic gradient methods for nested composition optimization. *SIAM Journal on Optimization*, 29(1):616–659, 2019.

[67] Christos Zacharias, Nan Liu, and Mehmet A Begen. Dynamic inter-day and intra-day scheduling. *Available at SSRN 3728077*, 2020.

[68] Christos Zacharias and Michael Pinedo. Managing customer arrivals in service systems with multiple identical servers. *Manufacturing & Service Operations Management*, 19(4):639–656, 2017.

[69] Zhe Zhang and Guanghui Lan. Optimal algorithms for convex nested stochastic composite optimization. *arXiv preprint arXiv:2011.10076*, 2020.

[70] Serhan Ziya. Queues with no-show customers. *Queueing Systems*, pages 1–3, 2022.

APPENDICES

Appendix A

Proofs

A.1 Proof of Theorem 1

By definition (4.2), $Pr\{e_{i\tau} \le r_i\}$ is the probability of rejecting at most r_i patients in class i. This probability is obviously greater than $Pr\{|e_{i\tau}| \le r_i\}$ since for any random variable Z and any nonnegative parameter z, we have $Pr\{|Z| \le z\} = Pr\{Z \le z\} - Pr\{Z \le -z\}$. Therefore,

$$\begin{aligned} Pr\{e_{i\tau} \le r_i\} &\ge Pr\{|e_{i\tau}| \le r_i\} \\ &= 1 - Pr\{|e_{i\tau}| \ge r_i\} \\ &\ge 1 - \frac{\mathbb{E}\left[e_{i\tau}^2\right]}{r_i^2} \\ &\ge 1 - \epsilon, \end{aligned}$$

where the second inequality follows from the Markov's inequality in second moments, and the last inequality follows because the expected least square error is bounded by m_i and if we define r_i as $r_i = \sqrt{\frac{m_i}{\epsilon}}$.

A.2 Proof of Theorem 2

Strong duality holds, if the primal problem is convex and Slater's condition holds (see e.g. [8]). By definition (3.6), $f(x_t)$ is convex in x_t as it is the sum of a convex overtime cost and linear waiting cost. Thus, the objective function of problem (4.1) is convex in X_t. Also, the feasible set is a convex set since all constraints are linear-quadratic functions, hence the primal problem is convex.

It remains to show the Slater's condition holds, i.e., there exists a strictly feasible solution. Recall the definitions of $\mathcal{C}_t(s_t)$ and $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$ in (3.4) and (4.3), respectively. Considering the capacity constraints, there must exist a strictly feasible solution for a large enough overtime usage independent of the scheduling decisions. Therefore, we just need to show that the Slater's condition holds for the demand constraints that are $\sum_{n=0}^{N-1} a_{i,n}^t = d_{i,t}, \forall i \in [I]$ in $\mathcal{C}_t(s_t)$ (deterministic demand constraints) and $\mathbb{E}\left[e_{i\tau}^2\right] \leq m_i, \forall i \in [I]$ in $\bar{\mathcal{C}}_\tau(\bar{s}_\tau)$ (stochastic demand constraints). We just need to find $a_{i,n}^t$ and $a_{i,n}^\tau$ such that the deterministic demand constraints holds and the stochastic demand constraints are strict. Since finding the strictly feasible solution is independent of scheduling decisions, it is obvious that there exist a strictly feasible solution for the deterministic demand constraints. conversely, it is not obvious for the stochastic constraints since we have

$$\mathbb{E}\left[e_{i\tau}^2\right] = \mathbb{E}\left[\left(D_{i,\tau} - \sum_{n=0}^{N-1} a_{i,n}^\tau\right)^2\right] = \mathrm{Var}(D_{i,\tau}) + \left(\mathbb{E}\left[D_{i,\tau}\right] - \sum_{n=0}^{N-1} a_{i,n}^\tau\right)^2 \leq m_i$$

for all classes of patients. However, since the variance of arrivals of patients is bounded, we can choose m_i such that $m_i > \max\{\mathrm{Var}(D_{i,t+1}), \ldots, \mathrm{Var}(D_{i,T})\}$, hence we can guarantee there exists a strictly feasible in stochastic demand constraint as well, hence the Slater's condition holds.

A.3 Proof of Lemma 1

It follows from Theorem 2 that $\mathcal{Y}_t(s_t)$ is a convex set. Therefore, it remains to show that for a given sample path $D_{]t,T]}$, the random loss function $\ell_t(X_t, D_{]t,T]})$ is a convex function in X_t since it implies that $\mathbb{E}_{D_{]t,T]}}\left[\ell_t(X_t, D_{]t,T]})\right]$ is also convex in X_t. We prove the loss function is convex by showing that the Hessian matrix of loss function is positive semi-definite, e.g. all eigenvalues are nonnegative. We start with finding the forms of second partial derivatives of the loss function for a given demand d. Recall $X_t = (x_t, \bar{x}_{]t,T]})$ includes both scheduling and allocation decisions, thus we have two groups of variables $a_{i,n}^\tau$ and o_τ where $\tau \in \{t, t+1, \ldots, T\}$. Then,

$$\frac{\partial^2 \ell_t(X_t, D_{]t,T]})}{\partial a_{i_1,n_1}^{\tau_1} \partial a_{i_2,n_2}^{\tau_2}} = \begin{cases} 0 & i_1 \neq i_2 \text{ or } \tau_1 \neq \tau_2 \\ 2\lambda_{i\tau}^t & \text{otherwise } (i_1 = i_2 = i \text{ and } \tau_1 = \tau_2 = \tau) \end{cases},$$

$$\frac{\partial^2 \ell_t(X_t, D_{]t,T]})}{\partial o_{\tau_1} \partial o_{\tau_2}} = \begin{cases} 0 & \tau_1 \neq \tau_2 \\ u''_\tau(o_\tau) & \tau_1 = \tau_2 = \tau \end{cases},$$

$$\frac{\partial^2 \ell_t(X_t, D_{]t,T]})}{\partial a_{i,n}^{\tau_1} \partial o_{\tau_2}} = 0,$$

where $u''_\tau(o_\tau) \geq 0$ since the overtime cost function is convex, and we define $\lambda_{it}^t = 0$ for all $i \in [I]$. Using the above results, the Hessian matrix is a *block diagonal* matrix with following form

$$H_{\ell_t} = \begin{bmatrix}
B_{1,t} & & & & \cdots & & & & 0\\
& \ddots & & & & & & & \\
& & B_{1,T} & & & & & & \\
& & & \ddots & & & & & \\
& & & & \ddots & & & & \\
\vdots & & & & B_{I,t} & & & & \vdots\\
& & & & & \ddots & & & \\
& & & & & & B_{I,T} & & \\
& & & & & & & u''_\tau(o_t) & \\
& & & & & & & & \ddots \\
0 & & & & \cdots & & & & u''_\tau(o_T)
\end{bmatrix}_{r_t \times r_t}$$

where $T' = T - t + 1$ and $B_{i,\tau} = 2\lambda_{i\tau}^t \mathbb{1}_N \mathbb{1}_N^\top$. Using the fact that the determinant of a block diagonal matrix is the product of the determinants of the diagonal blocks, we have

$$\det(H_{\ell_t} - \lambda \mathbb{I}_{r_t}) = \prod_{\tau=t}^{T} \left[(u''_\tau(o_\tau) - \lambda) \prod_{i=1}^{I} \det(B_{i,\tau} - \lambda \mathbb{I}_N) \right].$$

Now we need to find the eigenvalues of each matrix $B_{i,\tau}$. If $B_{i,\tau} v = \lambda v$ with $v \neq 0$, then $\lambda v = B_{i,\tau} v = 2\lambda_{i\tau}^t \langle \mathbb{1}_N, v \rangle \mathbb{1}_N$. If $\langle \mathbb{1}_N, v \rangle \neq 0$ which means that v is a multiple of $\mathbb{1}_N$, and the corresponding eigenvalue obtained for $v = \mathbb{1}_N$ that is $\lambda = 2N\lambda_{i\tau}^t$. If $\langle \mathbb{1}_N, v \rangle = 0$ which happens for $N-1$ linearly independent vectors, we have $\lambda = 0$. Consequently, we have

$$\det(H_{\ell_t} - \lambda \mathbb{I}_{r_t}) = \prod_{\tau=t}^{T} \left[(u''_\tau(o_\tau) - \lambda) \lambda^{I(N-1)} \prod_{i=1}^{I} (2N\lambda_{i\tau}^t - \lambda) \right].$$

Therefore, the eigenvalues of the Hessian matrix of the loss function have the values 0, $u''_\tau(o_\tau)$, and $2N\lambda^t_{i\tau}$ for all $i \in [I]$ and $\tau \in \{t, \dots, T\}$, and obviously, all of them are nonnegative.

A.4 Proof of Theorem 3

We start with the definition of AR in (4.9), and we prove it equals to the sum of the expectation of LRs. The general idea is to use two offline problem (4.10) and (4.11) to build the definition of LR. Based on (4.11), $V^1_2(s_1, d) = \text{OPT}(d)$ that $\text{OPT}(d)$ is the optimal solution of the offline problem defined in (3.7). Also, problem (4.10) implies $V^T_1(s_T, x^{alg}_T, d) = F(x^{alg})$ that $F(x^{alg})$ is the total cost of scheduling using Algorithm 1. Therefore, we have

$$
\begin{aligned}
\text{AR}^T_{alg} &= \mathbb{E}\left[F(x^{alg})\right] - \mathbb{E}\left[\text{OPT}(d)\right] \\
&= \mathbb{E}\left[F(x^{alg}) - V^1_2(s_1, d)\right] \\
&= \mathbb{E}\left[F(x^{alg}) - V^1_1(s_1, x^{alg}_1, d) + V^1_1(s_1, x^{alg}_1, d) - V^1_2(s_1, d)\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[V^1_1(s_1, x^{alg}_t, d) - V^1_2(s_1, d)|s_1\right]\right] + \mathbb{E}\left[F(x^{alg}) - V^1_1(s_1, x^{alg}_1, d)\right] \\
&= \mathbb{E}\left[\text{LR}_1(s_1, x^{alg}_t)\right] + \mathbb{E}\left[F(x^{alg}) - V^2_2(s_2, d)\right] = \dots \\
&= \sum_{t=1}^{T} \mathbb{E}\left[\text{LR}_t(s_t, x^{alg}_t)\right] + \mathbb{E}\left[F(x^{alg}) - V^T_1(s_T, x^{alg}_T, d)\right] \\
&= \sum_{t=1}^{T} \mathbb{E}\left[\text{LR}_t(s_t, x^{alg}_t)\right],
\end{aligned}
$$

where the fourth equality follows from the law of iterated expectations, and the fifth equality follows from the definition of LR in (4.12) and the fact that the arrivals are known in offline problem, thus the only difference between s_t and s_{t-1} is in x^{alg}_{t-1} that is given.

A.5 Proof of Lemma 2

It follows from Theorem 2 that $\mathcal{Y}_t(s_t)$ is a convex set. Therefore, it remains to show that the objective function in (5.6) is convex. By definition (5.3), the risk measure $\rho(.)$ is coherent, thus it is convex and monotone. By Lemma 1, a given sample path $D_{]t,T]}$, the random loss function $\ell_t(X_t, D_{]t,T]})$ is a convex function in X_t. The loss function is convex and the risk measure is convex and non-decreasing, then $\rho(\ell_t(X_t, D_{]t,T]}))$ is convex in X_t.

A.6 Proof of Lemma 3

Lipschitz Smoothness: We prove all layer functions $h_{1,t}, h_{2,t}, h_{3,t}$ and their derivative are Lipschitz continuous.

- $h_{3,t}(X_t) = \mathbb{E}\left[\ell_t(X_t, D_{]t,T]})\right]$. Based on the Mean Value Theorem, it is enough to show that the first and second order partial derivatives are bounded. The general format of each element of $\nabla h_{3,t}$ can have two below forms

$$\frac{\partial h_{3,t}(X_t)}{\partial a^{\tau}_{i,n}} = \mathbb{E}\left[\frac{\partial \ell(X, D_{]t,T]})}{\partial a^{\tau}_{i,n}}\right] = \left\{\begin{array}{cc} n\gamma^w_i & \tau = t \\ n\gamma^w_i - 2\lambda^t_{i\tau}\left(\mathbb{E}\left[D_{i,\tau}\right] - \sum_{n'=0}^{N} a^{\tau}_{i,n'}\right) & t < \tau \leq T \end{array}\right., \quad \text{(A.1)}$$

$$\frac{\partial h_{3,t}(X_t)}{\partial o_\tau} = \mathbb{E}\left[\frac{\partial \ell(X_t, D_{]t,T]})}{\partial o_\tau}\right] = u'_\tau(o_\tau), \quad \text{(A.2)}$$

  where both forms are bounded since scheduling decisions and overtime are bounded. Also, the general format of each element of $\nabla^2 h_{3,t}$ can have two below forms

$$\frac{\partial^2 h_{3,t}(X_t)}{\partial (a^{\tau}_{i,n})^2} = \left\{\begin{array}{cc} 0 & \tau = t \\ 2\lambda^t_{i\tau} & t < \tau \leq T \end{array}\right.,$$

$$\frac{\partial^2 h_{3,t}(X_t)}{\partial (o_\tau)^2} = u''_\tau(o_\tau),$$

  where both are bounded since the overtime is bounded and the overtime cost is a smooth convex function of overtime. Hence $h_{3,t}$ and its derivative are Lipschitz continuous.

- $h_{2,t}(X_t, \omega_3) = \mathbb{E}\left[\left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]^2_+\right]$. We define $g(X_t, \omega_3) = \left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]^2_+$, and aim to show that $g(X_t, \omega_3)$ and and its derivative are Lipschitz continuous for a given sample path $D_{]t,T]}$ that easily implies $h_{2,t}(X_t, \omega_3)$ and its derivative are Lipschitz continuous. Before that, we establish Lemma 4 as follows:

Lemma 4 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a smooth function. We define $x_1, x_2 \in \mathbb{R}^n$ and $u_1, u_2 \in \mathbb{R}$, such that $f(x_1) > u_1$ and $f(x_2) < u_2$. Then, there exists at least a parameter $\beta \in [0,1]$ such that $f(\beta x_1 + (1-\beta)x_2) = \beta u_1 + (1-\beta)u_2$.*

Proof: We define $q : \mathbb{R} \to \mathbb{R}$ and $d : \mathbb{R} \to \mathbb{R}$ as follows

$$q(t) = f\big(tx_1 + (1-t)x_2\big).$$
$$d(t) = tu_1 + (1-t)u_2.$$

Clearly, $d(t)$ is a linear function. Since f is smooth, $q(t)$ is also a smooth function. Since we know that $q(1) > d(1)$ and $q(0) < d(0)$, there should be at least one t' such that $q(t') = d(t')$. See Figure A.1 for both states $u_2 < u_1$ and $u_2 > u_1$. □

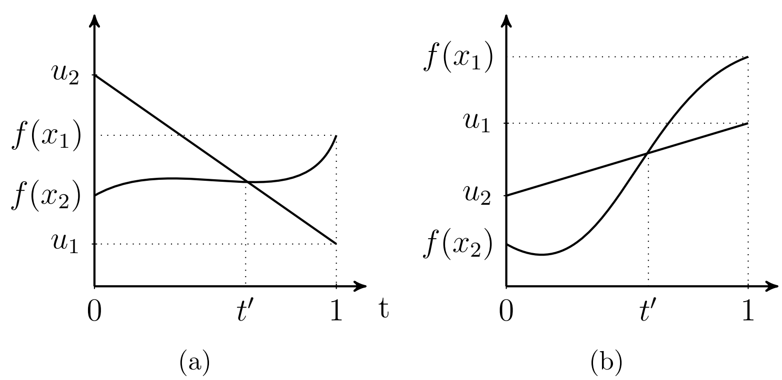

Figure A.1: $q(t)$ vs $d(t)$ (a) $u_2 > u_1$ (b) $u_2 < u_1$

The general format of each element of ∇g can have three below forms

$$\frac{\partial g(X_t,\omega_3)}{\partial a^{\tau}_{i,n}} = \begin{cases} 2n\gamma^w_i.\left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]_+ & \tau = t \\ 2\left(n\gamma^w_i - 2\lambda^t_{i\tau}\left(\mathbb{E}\left[D_{i,\tau}\right] - \sum_{n'=0}^{N} a^{\tau}_{i,n'}\right)\right).\left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]_+ & t < \tau \leq T \end{cases},$$

$$\frac{\partial g(X_t,\omega_3)}{\partial o_\tau} = 2u'_\tau(o_\tau).\left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]_+,$$

$$\frac{\partial g(X_t,\omega_3)}{\partial \omega_3} = -2\left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]_+.$$

Again, patients arrivals and the decision variables are bounded, and then the loss function is bounded. Therefore, all above forms are bounded. Then, $g(X_t, \omega_3)$ is Lipschitz continuous for a given sample path $D_{]t,T]}$. It remains to show that the gradient of $g(X_t, \omega_3)$ is also Lipschitz continuous. We define $q(X_t, \omega_3) = \ell(X_t, D_{]t,T]}) - \omega_3$. Clearly, $\nabla g(X_t, \omega_3)$ is not differentiable for any X_t and ω_3 such that $q(X_t, \omega_3) = 0$. Therefore, Mean Value theorem cannot be used in this case. We divide our discussion into three states. Then, for all $X_t^1, X_t^2 \in \mathbb{R}^{r_t}$ and $\omega_3^1, \omega_3^2 \in \mathbb{R}$:

1. if $q(X_t^1, \omega_3^1) \leq 0$ and $q(X_t^2, \omega_3^2) \leq 0$, then $g(X_t^1, \omega_3^1) = g(X_t^2, \omega_3^2) = 0$; and therefore
$$|g(X_t^1, \omega_3^1) - g(X_t^2, \omega_3^2)| = 0.$$
2. if $q(X_t^1, \omega_3^1) \geq 0$ and $q(X_t^2, \omega_3^2) \geq 0$, then $g(X_t, \omega_3) = q^2(X_t, \omega_3)$ and it is enough to show that the gradient of $q^2(X_t, \omega_3)$ is Lipschitz continuous for all $X_t \in \mathbb{R}^r$ and $\omega_3 \in \mathbb{R}$. The general form of each element of $\nabla^2 q^2$ can have three below forms
$$\frac{\partial^2 q^2(X_t, \omega_3)}{\partial (a_{i,n}^\tau)^2} = 2\left(n\gamma_i^w - 2\lambda_{i\tau}^t\Big(\mathbb{E}\left[D_{i,\tau}\right] - \sum_{n'=0}^{N} a_{i,n'}^\tau\Big)\right)^2 + 2\lambda_{i\tau}^t\Big(\ell_t(X_t, D_{]t,T]}) - \omega_3\Big),$$
$$\frac{\partial^2 q^2(X_t, \omega_3)}{\partial (o_\tau)^2} = 2u_\tau''(o_\tau).q(X_t, \omega_3) + 2\left(u_\tau'(o_\tau)\right)^2,$$
$$\frac{\partial q^2(X_t, \omega_3)}{\partial \omega_3} = 2.$$
From previous layer, ω_3 is bounded since $\mathbb{E}\left[\ell_t(X_t, D_{]t,T]})\right]$ is bounded. Also, patients arrivals and the decision variables are bounded, and then all above forms are bounded. Hence the Mean Value Theorem implies ∇q^2 is Lipschitz continuous.
3. if $q(X_t^1, \omega_3^1).q(X_t^2, \omega_3^2) < 0$, then $g(X_t^1, \omega_3^1) = (\ell_t(X_t^1, D_{]t,T]}) - \omega_3^1)^2$ and $g(X_t^2, \omega_3^2) = 0$. Lemma 4 states that there exists a pair (X_t^3, ω_3^3) such that $q(X_t^3, \omega_3^3) = 0$ and
$$\left\|(X_t^1, \omega_3^1) - (X_t^3, \omega_3^3)\right\| \leq \left\|(X_t^1, \omega_3^1) - (X_t^2, \omega_3^2)\right\|. \tag{A.3}$$
From state 2 we know that there is a L such that
$$|g(X_t^1, \omega_3^1) - g(X_t^3, \omega_3^3)| \leq L\left\|(X_t^1, \omega_3^1) - (X_t^3, \omega_3^3)\right\|,$$
and also from (A.3) and the fact that $g(X_t^3, \omega_3^3) = g(X_t^2, \omega_3^2) = 0$ we have
$$|g(X_t^1, \omega_3^1) - g(X_t^2, \omega_2)| \leq L\left\|(X_t^1, \omega_3^1) - (X_t^2, \omega_3^2)\right\|.$$

Therefore, $g(X_t, \omega_3)$ and and its derivative are Lipschitz continuous for a given sample path $D_{]t,T]}$. Then, $h_{2,t}(X_t, \omega_3)$ and its derivative are Lipschitz continuous. Note that the Lipschitz constant would be the maximum constant obtained in states 1 to 3. □

- $h_{1,t}(X_t, \omega_2) = \mathbb{E}\left[\ell_t(X_t, D_{]t,T]}) + \varkappa\sqrt{\varepsilon^2 + \omega_2}\right]$. We first show that $\sigma(\omega_2) = \sqrt{\varepsilon^2 + \omega_2}$ and its derivative are Lipschitz continuous for all $\omega_2 \in \mathbb{R}^+$. We have
$$\frac{\partial \sigma(\omega_2)}{\partial \omega_2} = \frac{1}{2\sqrt{\epsilon^2 + \omega_2}} \leq \frac{1}{2\epsilon},$$

$$\frac{\partial^2 \sigma(\omega_2)}{\partial \omega_2^2} = \frac{1}{4(\epsilon^2+\omega_2)^{\frac{3}{2}}} \leq \frac{1}{4\epsilon^3}.$$

that yields $\sigma(\omega_2)$ and its derivative are Lipschitz continuous according to Mean Value Theorem. Now we show that $h_{1,t}(X_t, \omega_2)$ is Lipschitz continuous since for all $X_t^1, X_t^2 \in \mathbb{R}^{r_t}$ and $\omega_2^1, \omega_2^2 \in \mathbb{R}$, we have

$$\begin{aligned}
\left|h_{1,t}(X_t^1, \omega_2^1) - h_{1,t}(X_t^2, \omega_2^2)\right| &= \left|\mathbb{E}\left[\ell(X_t^1, D_{]t,T]}) + \varkappa\sqrt{\varepsilon^2+\omega_2^1}\right]\right. \\
&\qquad \left. - \mathbb{E}\left[\ell(X_t^2, D_{]t,T]}) + \varkappa\sqrt{\varepsilon^2+\omega_2^2}\right]\right| \\
&\leq \left|\mathbb{E}\left[\ell(X_t^1, D_{]t,T]})\right] - \mathbb{E}\left[\ell(X_t^2, D_{]t,T]})\right]\right| \\
&\qquad + \varkappa\left|\sqrt{\varepsilon^2+\omega_2^1} - \sqrt{\varepsilon^2+\omega_2^2}\right| \\
&\leq \left|h_{3,t}(X_t^1) - h_{3,t}(X_t^2)]\right| + |\sigma(\omega_2^1) - \sigma(\omega_2^2)| \\
&\leq L_{h_{3,t}}\left\|X_t^1 - X_t^2\right\| + L_\sigma|\omega_2^1 - \omega_2^2|,
\end{aligned}$$

that the last inequality follows from the Lipschitz continuity of $h_{3,t}$ and σ, and note that $\varkappa \in [0,1]$; then, $h_{1,t}$ is Lipschitz continuous. Similarly, the derivative of $h_{1,t}$ is also Lipschitz continuous since the derivatives of $h_{3,t}$ and σ are Lipschitz continuous.

□

Appendix B

Technical Details on SDSA

As discussed in Chapter 4, relaxing the anticipativity constraints simplifies the high-dimension MAPS problem where we only need to provide an algorithm to solve problem (4.1) every day as we move forward to the end of the time horizon. Theorem 2 shows that the subgradient method can guarantee both convergence and optimality due to the strong duality of model (4.5). However, the optimality is guaranteed if we can find the optimal solution of the sub-problem (4.7) in each iteration. In Algorithm 2, we present a subgradient method in which we optimally solve the stochastic problem (4.7) in each iteration. We next provide a detailed explanation on Steps 1-4 in the following sections.

B.1 Illustration for Step 1 in SDSA

In Step 1, we employ an unified optimal method, namely, the accelerated stochastic approximation algorithm (proposed by [15]) to optimally solve the stochastic sub-problem (4.7). In particular, this problem is a constrained convex stochastic optimization, and this method is a modified version of Nesterov's algorithm (see [44] and [43]).

In Step 1.1., We need two arbitrarily feasible solutions (that can be the same) for this step, and since all constraints are linear, it is easy to find these solutions. We also need to define three parameters where we need to compute the parameter γ in Step 1.2. Let M_k be the Lipschitz constant of $\mathbb{E}_{D_{]t,T]}}\left[L_t(X_t, D_{]t,T]}, \lambda_k^t)\right]$, and V be $\frac{1}{2}\left\|Y_0 - X_{t,k}\right\|^2$, and σ^2 be an upper bound for the expected square norm of estimated stochastic subgradient and the actual subgradient (more details in [15]).

Algorithm 2: Single-Day Stochastic Algorithm (SDSA)

Input: Number of iterations (R), Sample pool($\mathcal{D}$).

Initial: Set $k = 0$, and choose $\lambda_0^t \in \mathbb{R}^{I(T-t)}$.

while *True* **do**

1. To solve $X_{t,k} = \min_{X_t \in \mathcal{Y}_t} \quad \mathbb{E}_{D_{]t,T]}} \left[L_t(X_t, D_{]t,T]}, \lambda_k^t) \right]$, **do**

1.1. Choose $Y_0^{alg} \in \mathcal{Y}_t$ and $Y_0 \in \mathcal{Y}_t$, and compute M_k, V, and σ.

1.2. Set $\gamma = \sqrt{\frac{(\sigma^2+4M_k^2)R(R+1)(R+2)}{3V}}$.

1.3. **for** $r = 0, 1, \ldots, R-1$ **do**

1.3.1. Set $\alpha_r = \frac{2}{r+2}$ and $\gamma_r = \frac{4\gamma}{(r+1)(r+2)}$.

1.3.2. Compute $Y_{r+1}^{md} = (1 - \alpha_r) Y_r^{ag} + \alpha_r Y_r$.

1.3.3. Randomly pick the sample D_r from the sample pool $\mathcal{D}$.

1.3.4. Use the sample D_r to compute the stochastic gradient G_r in Y_{r+1}^{md}.

1.3.5. Solve $Y_{r+1} = \underset{Y \in \mathcal{Y}_t}{\operatorname{argmin}} \left\{ \alpha_r \langle G_r, Y \rangle + \frac{\gamma_r}{2} \|Y - Y_r\|^2 \right\}$.

1.3.6. Compute $Y_{r+1}^{ag} = \alpha_r Y_{r+1} + (1 - \alpha_r) Y_r^{ag}$.

1.4. Set $X_{t,k} = Y_R^{ag}$.

2. Set $\bar{L}_{\lambda_k^t} = \mathbb{E}_{D_{]t,T]}} \left[L_t(X_{t,k}, D_{]t,T]}, \lambda_k^t) \right]$, and choose subgradient $g_k = \mathbb{E} \left[e_{i\tau}^2 - m_i \right]$.

3. Compute $\lambda_{k+1}^t = [\lambda_k^t + \theta_k g_k]_+$ where θ_k is the stepsize.

4. **If** $\left\| \lambda_{k+1}^t - \lambda_k^t \right\| = 0$, **then** *break* the loop, **else** set k = k + 1 and go to Step 1.

Output: Lagrangian multiplier $\lambda^t = \lambda_k^t$ and scheduling decision $X_t = X_{t,k}$.

As we mentioned, our algorithm is data-riven in the sense that in Step 1.3.3, we use the offline data set to randomly pick a sample arrival for each estimation of the algorithm. Also, it is worthy to mention that the optimization problem in Step 1.3.5 is a quadratic optimization that can be solved by the state-of-art solvers (e.g. Gurobi) very fast. In fact, Y_{r+1} is the projection of $Y_r - \alpha_r G_r/\gamma_r$ into $\mathcal{Y}_t$ in each iteration.

B.2 Illustration for Step 2-4 in SDSA

In Step 2 and Step 3, we need to find the subgradient and the stepsize to update the Lagrangian multiplier for the next iteration. In fact, we use the relaxed demand constraint to compute the subgradient. Then, we set the Polyak stepsize as $\theta_k = \frac{UB - \bar{L}_{\lambda_k^t}}{\|g_k\|^2}$. We should also ensure that all elements of λ_k^t are nonnegative. Note that in each iteration we use Stochastic Approximation (SA) method to approximate both subgradient and $\bar{L}_{\lambda_k^t}$. Finally, we need to set a small number δ to check when $\left\|\lambda_{k+1}^t - \lambda_k^t\right\| \leq \delta$ to stop the subgradient algorithm.

B.3 Numerical Results on the Convergence of SDSA

In this section, we provide some numerical results to indicate the convergence of Algorithm 2. To this end, Figure B.1 shows $\left\|\lambda_{k+1}^t - \lambda_k^t\right\|$ converges to zero, and also it indicates that $\bar{L}_{\lambda_k^t}$ converges to a fixed number which is the (near-)optimal solution. Note that the numerical setting and all parameters are presented in Table 6.2.

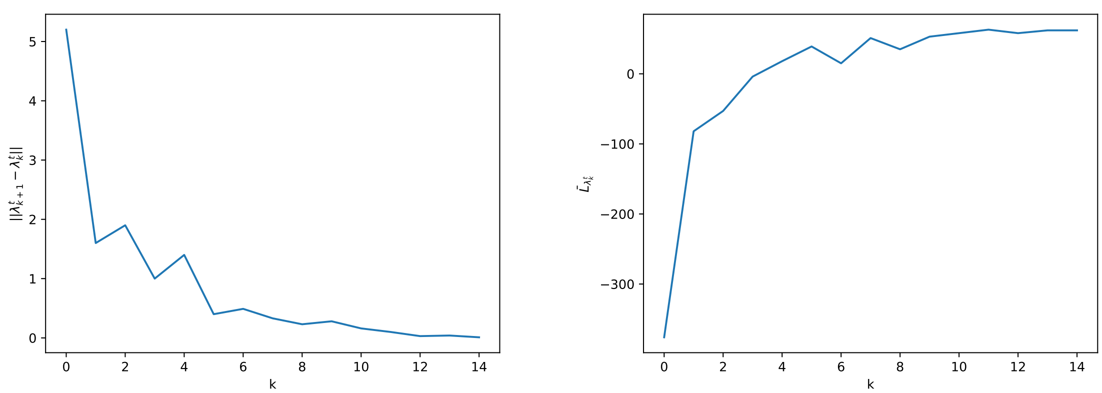

(a) (b)

Figure B.1: Convergence of SDSA (a) Lagrangian multiplier (b) Lagrangian function

Appendix C

Technical Details on SDRA

In this section, we provide technical details required for the SDRA in which we implement a multi-level nested averaging stochastic gradient method from [6]. The outline of SDRA is presented in Algorithm 3. [6] shows that this algorithm converges to a stationary point in nonconvex problems. However, according to Lemma 2, our problem is convex meaning that the proposed algorithm converges to the optimal point. We next provide a detailed explanation on Steps 1-4 in the following sections.

C.1 Illustration for Step 1 in SDRA

In Step 1, we first compute U^k by solving a quadratic optimization. Similar to our discussion in Step 1.3.5 in SDSA, U^k is the projection of $X^{t,k} - \frac{\mathcal{Z}^k}{2}$ onto $\mathcal{Y}_t$. We utilize the solution of this quadratic optimization in order to update $X^{t,k}$. Note that we set the value of the step size τ_k and the parameter β based on the results in Theorem 2.1 in [6].

C.2 Illustration for Step 2 in SDRA

In Step 2, we need to obtain some statistical estimates. To do so, we pick six independent sample arrivals and utilize them in order to compute the value functions and the stochastic gradients. We provide a detailed formulation to compute these statistical estimates.

Algorithm 3: Single-Day Robust Algorithm (SDRA)

Input: Number of iterations (K), Sample pool($\mathcal{D}$).

Initial: Obtain the optimal Lagrangian multiplier λ_t by running SDSA(s_t). Select an arbitrary $X_{t,0} \in \mathcal{Y}_t$. Let $W_1^0 = 0$, and $W_2^0 = W_3^0 = [\mathbf{0}]_{r+1}$. Set $\beta = 2$. Define the integer sequences $\{\tau_k\}_{k\geq 0}$ such that $\tau_0 = 1$ and $\tau_k = \frac{1}{\sqrt{K}}$ for all $k = 1, \cdots, K-1$.

for $k = 0, 1, \ldots, K-1$ **do**

1. Compute $U^k = \underset{Y\in\mathcal{Y}_t}{\operatorname{argmin}}\left\{\langle \mathcal{Z}^k, Y - X_{t,k}\rangle + \frac{\beta}{2}\left\|Y - X_{t,k}\right\|^2\right\}$, and set

$$X_{t,k+1} = (1-\tau_k)X_{t,k} + \tau_k U^k. \tag{C.1}$$

2. Obtain stochastic gradients $J_{1,t}^{k+1}, J_{2,t}^{k+1}, J_{3,t}^{k+1}$ and function values $G_{1,t}^{k+1}, G_{2,t}^{k+1}$, $G_{3,t}^{k+1}$ using six independent sample arrivals and the inner function estimates W_2^{k+1} and W_3^{k+1}.

3. Construct a biased estimate of a subgradient of the composite function:

$$g_3^{k+1} = J_{3,t}^{k+1}, \tag{C.2}$$

$$g_2^{k+1} = J_{2X_t,t}^{k+1} + J_{2\omega_3,t}^{k+1} g_3^{k+1}, \tag{C.3}$$

$$g_1^{k+1} = J_{1X_t,t}^{k+1} + J_{1\omega_2,t}^{k+1} g_2^{k+1}. \tag{C.4}$$

4. Update the path averages:

$$\mathcal{Z}^{k+1} = (1-\tau_k)\mathcal{Z}^k + \tau_k[g_1^{k+1}]^\top, \tag{C.5}$$

$$W_i^{k+1} = (1-\tau_k)W_i^k + \frac{\tau_k}{\lfloor\sqrt{K}\rfloor}\sum_{j=1}^{\lfloor\sqrt{K}\rfloor} G_{i,j}^{k+1}, \quad i \in \{2,3\}. \tag{C.6}$$

Output: Scheduling decision $X_t = X_{t,K}$.

Computing the Value Functions. According to the layer functions defined in (5.8), the value functions G_i for $i = \{1, 2, 3\}$ can be represented as follows,

$$G_{1,t}(X_t, \omega_2, D_{]t,T]}) = \ell_t(X_t, D_{]t,T]}) + \varkappa\sqrt{\varepsilon^2 + \omega_2}, \tag{C.7a}$$

$$G_{2,t}(X_t, \omega_3, D_{]t,T]}) = \left[\ell_t(X_t, D_{]t,T]}) - \omega_3\right]_+^2, \tag{C.7b}$$

$$G_{3,t}(X_t, D_{]t,T]}) = \ell_t(X_t, D_{]t,T]}), \tag{C.7c}$$

where

$$h_{1,t}(X_t, \omega_2) = \mathbb{E}\left[G_{t,1}(X_t, \omega_2, D_{]t,T]})\right], \tag{C.8a}$$

$$h_{2,t}(X_t, \omega_3) = \mathbb{E}\left[G_{t,2}(X_t, \omega_3, D_{]t,T]})\right], \tag{C.8b}$$

$$h_{3,t}(X_t) = \mathbb{E}\left[G_{t,3}(X_t, D_{]t,T]})\right]. \tag{C.8c}$$

Computing the Stochastic Gradients. The function $\ell_t(X_t, D_{]t,T]})$ is a smooth convex function. Then, it is differentiable with respect to X_t and integrable with respect to $D_{]t,T]}$. Thus, the expected value of the loss function that is used in (5.8c) is also differentiable. Although its derivative is not readily available, we can draw a sample $\tilde{\mathcal{D}}$ from the sample pool $\mathcal{D}$ and use an element of $\partial_{X_t}\ell_t(X_t, \tilde{\mathcal{D}})$ as a *stochastic subgradient* (a random vector whose expected value is a subgradient). Let J_1, J_2 and J_3 be the stochastic subgradient of $h_{1,t}$, $h_{2,t}$, and $h_{3,t}$ (and the subgradient of $G_{1,t}, G_{2,t}$, and $G_{3,t}$) generated from sample $\tilde{D}$, respectively. Then, we define $J_{1,t}^{k+1}$ and $J_{2,t}^{k+1}$ as follows

$$J_{i,t} = \begin{bmatrix} J_{iX_t,t} \\ J_{i\omega_{i+1},t} \end{bmatrix}_{(r+1)\times 1},$$

where $J_{iX_t,t}$ and $J_{i\omega_{i+1},t}$ are the stochastic subgradients of $h_{i,t}$ with respect to X_t and ω_2, respectively, for $i = 1, 2$. Furthermore, we define $J_{3,t}^{k+1}$ that is the stochastic gradient of the inner layer:

$$J_{3,t} = \left[J_{3X_t,t}\right]_{r\times 1},$$

that $J_{3X_t,t}$ is the stochastic subgradient of $h_{3,t}$ with respect to X_t.

C.3 Illustration for Step 3-4 in SDRA

In Step 3, we use the stochastic subgradients that were computed in Step 2 to construct a biased estimate of a subgradient of the composite function. Finally, we utilize the subgradient of the composite function to update the path averages by backward recursion in (C.5). Furthermore, we require a batch of samples to update the path-averaged inner functions estimates. The number of samples in each batch is fixed to the rounded square root of the number of iterations as mentioned in [6].

C.4 Numerical Results on the Convergence of SDRA

In this section, we verify the convergence of Algorithm 3 by providing some numerical results. Let X^{ag}, Z^{ag}, U_{ag} be generated by the SDRA. Similar to [6], we define

$$V(X^{ag}, Z^{ag}) = \|U^{ag} - X^{ag}\|^2 + \|Z^{ag} - \nabla H(X_t)\|^2, \tag{C.9}$$

to indicate the convergence of the SDRA. $V(X^{ag}, Z^{ag})$ is an upper bound for the norm square of the gradient of the problem (see Definition 2.1 in [6]). We use the stochastic approximation method to estimate $\mathbb{E}\left[\sqrt{V(X^{ag}, Z^{ag})}\right]$. Figure C.1 shows $V(X^{ag}, Z^{ag})$ converges to zero. Note that the numerical setting and all parameters are presented in Table 6.2.

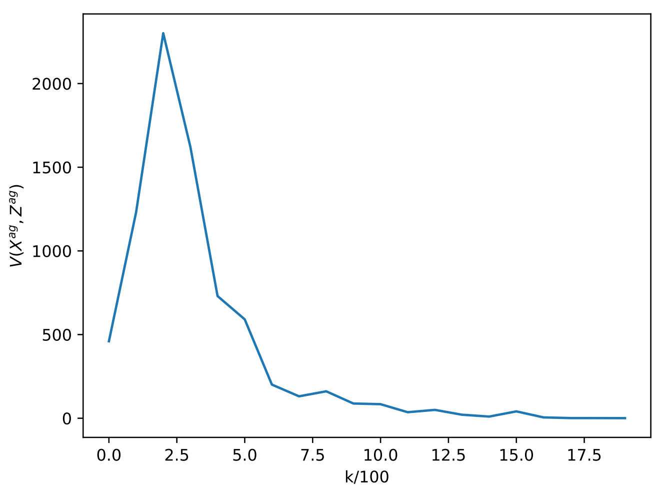

Figure C.1: Convergence of SDRA - $V(X^{ag}, Z^{ag})$